VSPAM.ORG · INDEPENDENT RESEARCH

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

A monthly research brief on the email, identity, and DNS threat landscape. July 2026 edition — the quarter's retrospectives land, and we score our own forecasts against them.

ISSUE	**Volume 2026, Issue 7**
ISSN	**2026-VSPAM-007**
REPORTING PERIOD	1 – 31 July 2026
RELEASED	August 2026
PUBLISHER	vSpam.org Independent Research
LICENSE	Creative Commons Attribution 4.0 (CC BY 4.0)

ABSTRACT

July 2026 delivered the Q2 retrospectives, and they resolved the year's central question. Microsoft's quarterly telemetry (23 July) showed Tycoon2FA-linked phishing down **92%** to 1.2 million messages — roughly **8% of its H2 2025 baseline** — with QR-code phishing falling from 18.7 million in March to **8.3 million** in June. Infrastructure disruption works. But Cisco Talos (28 July) recorded phishing in **over half** of all incident-response engagements, up from a third, and authentication abuse in **65%**, nearly double the prior quarter. Ransomware reached a 2026 high of **811 victims**. Check Point's Q2 brand data placed **ChatGPT in the top ten impersonated brands for the first time**. This issue scores the forecasts made in our May and June editions — including the one we got wrong.

Table of Contents

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

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

July was a month of answers. Three quarterly retrospectives published within six days — Microsoft's Q2 email threat landscape on 23 July, Check Point's Q2 brand phishing report the same day, and Cisco Talos' Q2 incident-response trends on 28 July — and between them they settled the question our May and June issues left open: does infrastructure disruption actually change the threat landscape?

The answer is yes, decisively, and with a qualification. Microsoft's telemetry shows Tycoon2FA-linked phishing at **1.2 million messages in June, a 92% reduction** from pre-disruption levels and approximately **8% of the 15.1 million monthly average recorded across H2 2025**.[1] The platform's share of CAPTCHA-gated phishing fell from 41% in March to 12% in June. QR-code phishing — the technique most closely associated with the platform — declined every month of the quarter, from a March peak of **18.7 million to 8.3 million** in June.

The qualification is that none of this made organisations safer. Cisco Talos found phishing accounting for **more than half of all incident-response engagements** in Q2, up from roughly a third in Q1, and authentication abuse present in **65% of engagements** against 35% the previous quarter.[2] Ransomware reached its 2026 high at **811 victims**.[3] Aggregate campaign volume fell; successful intrusion did not.

Four further signals define the month:

1. **ChatGPT entered the ten most-impersonated brands for the first time.** Microsoft led Check Point's Q2 ranking at **22.6%**, followed by LinkedIn (11.6%), Google (6.7%), Apple (5.8%), and Amazon (5.2%) — together more than half of all brand phishing. The ChatGPT campaigns used fake subscription-payment-failure notices to harvest full card details.[4]
2. **Ransomware hit a 2026 high with its top two operations tied.** 811 victims across 66 groups, up 15% from June. TheGentlemen and Qilin finished level at **119 victims each**, while June's debutant DeadLock collapsed from 81 to **22**.[3]
3. **Voice and collaboration channels absorbed the displaced volume.** Weekly malicious call attempts over Microsoft Teams reached **nearly ten times the mid-2025 baseline** by the end of Q2, with vishing attempts rising 31% and 27% in successive months.[1]
4. **In our own corpus, the fake-update technique outlived its operator.** A new ClickFix-family label entered our leading families in July even as ClearFake continued declining — direct evidence on the question our June issue committed to tracking. Section 6.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

The Takedown, Measured — Tycoon2FA and QR-Code Phishing Through Q2 2026

Platform-linked volume fell to 8% of its H2 2025 baseline; the associated technique declined every month

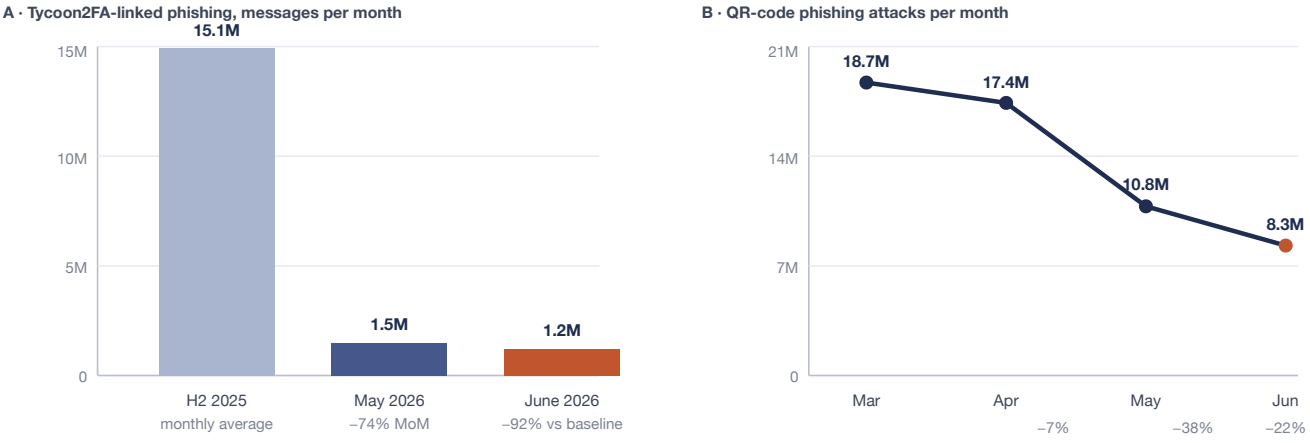

Source: Microsoft Security, "Email threat landscape: Q2 2026 trends and insights" (23 July 2026).

FIGURE 1 *The clearest evidence in years that infrastructure disruption moves aggregate volume. The platform seized in March fell to roughly a twelfth of its prior scale, and the technique it popularised declined in every subsequent month.*

Figure 1 is the empirical answer to a question this series has carried since March. It should be read alongside Section 2, which shows that the same quarter produced record levels of successful intrusion — because the two findings together are the actual lesson, and either one alone is misleading.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

Headline Numbers at a Glance

METRIC	VALUE	SOURCE
Email phishing threats detected, Q2 2026	**≈7.6 billion (Q1: 8.3B)**	Microsoft
Monthly trajectory, April to June	**2.7B → 2.4B**	Microsoft
Tycoon2FA-linked phishing, June	**1.2M (–92% vs baseline)**	Microsoft
Share of H2 2025 monthly average	**≈8% of 15.1M**	Microsoft
Platform share of CAPTCHA-gated phishing	**41% (Mar) → 12% (Jun)**	Microsoft
QR-code phishing, March to June	**18.7M → 8.3M (–56%)**	Microsoft
CAPTCHA-gated phishing, March to June	**12M → 2.2M (–81%)**	Microsoft
Credential phishing share of payload attacks	**94–96%**	Microsoft
ICS calendar-file payloads, June	**+277% MoM**	Microsoft
SVG share of CAPTCHA-gated payloads	**5% (Apr) → 26% (Jun)**	Microsoft
BEC attacks, April / May / June	**9M / 3.4M / 3.9M**	Microsoft
Teams malicious call attempts, end Q2	**≈10× mid-2025 baseline**	Microsoft
Phishing share of IR engagements, Q2	**>50% (Q1: ≈33%)**	Cisco Talos
Authentication abuse in engagements	**65% (Q1: 35%)**	Cisco Talos
Engagements with insufficient logging	**42% (Q1: 18%)**	Cisco Talos
Ransomware share of IR engagements	**>20%**	Cisco Talos
Most-targeted sector, Q2 IR	**Healthcare 17%**	Cisco Talos
Most-impersonated brand, Q2	**Microsoft, 22.6%**	Check Point
New entrant to top-ten impersonated brands	**ChatGPT (first time)**	Check Point
Ransomware victims posted to leak sites	**811 (66 groups, 78 countries)**	BreachSense
Month-on-month ransomware change	**+15% — 2026 high**	BreachSense
Joint most active groups	**TheGentlemen 119, Qilin 119**	BreachSense
Healthcare ransomware victims	**71 — highest since February**	BreachSense
Domains with no effective DMARC protection	**68.7%**	DmarcDkim
vSpam.org collector indicators ingested	**24,254**	vSpam.org
vSpam.org ThreatFox ingestion change	**1,084 (5.1× June)**	vSpam.org

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

REFERENCE

Methodology and Data Sources

This report consolidates public, attributable data published during or directly referencing July 2026, and 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 for this issue include Microsoft Security's quarterly email threat landscape reporting, Cisco Talos incident-response trend data, Check Point Research brand phishing analysis, BreachSense ransomware tracking, Hornetsecurity monthly threat reporting, CISA, and continuous DMARC measurement corpora.

A note on what this issue can now answer

Our May and June issues were written before the Q2 retrospectives published, and both flagged open questions accordingly. Those datasets are now available, and Section 7 scores every forecast this series has made against them — including one that was substantially wrong. We regard published, checkable forecasts as a basic obligation of a research series that expects to be taken seriously, and scoring them honestly as the other half of that obligation.

Quarter labels and monthly attribution

Microsoft's Q2 2026 reporting covers April to June 2026 and was published on 23 July. Cisco Talos' Q2 report covers the same period and was published on 28 July. Check Point's Q2 brand data covers April to June and published on 23 July. **None of these datasets describes July activity.** Where this issue reports July-specific figures they come from BreachSense's monthly series, Hornetsecurity's July digest, CISA's July catalogue additions, or our own corpus — and are labelled as such.

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.

As in previous issues, **community submission volume was low through this reporting window**, so no community-derived volume claims are made. Malware-family labels are inherited from upstream feed taxonomies which mix true family names with file architectures and behavioural tags; we identify which is which wherever it affects the analysis. All dates are normalised to UTC.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

SECTION 1

The Q2 Email Threat Landscape

1.1 Aggregate Volume

Microsoft detected approximately **7.6 billion email-based phishing threats** across Q2 2026, against 8.3 billion in Q1 — with monthly volumes declining through the quarter from **2.7 billion in April to 2.4 billion in June**.[1]

An 8% quarter-on-quarter decline in a metric that has risen consistently for years is a genuine inflection, and Section 1.2 attributes most of it. It is also worth keeping in proportion: 7.6 billion threats in a quarter is roughly 84 million per day. The landscape did not become quiet; one large contributor to it was removed.

1.2 The Takedown Verified

The March seizure of Tycoon2FA infrastructure produced the largest measurable effect on aggregate phishing volume this series has recorded.

MEASURE	VALUE	COMPARISON
Tycoon2FA-linked phishing, June 2026	**1.2M messages**	−92% from pre-disruption levels
H2 2025 monthly average	**15.1M messages**	June is ≈8% of this baseline
May 2026	**1.5M messages**	−74% month on month
Share of CAPTCHA-gated phishing, March	**41%**	—
Share of CAPTCHA-gated phishing, June	**12%**	−29 percentage points
QR campaigns redirecting to platform domains	**20% → 14%**	March to June

The associated techniques declined in step. **QR-code phishing** fell from a March peak of 18.7 million to 17.4 million in April (−7%), 10.8 million in May (−38%), and **8.3 million in June** (−22%) — a 56% reduction across the quarter. **CAPTCHA-gated phishing** fell further and faster, from a March peak of 12 million to **2.2 million in June**, an **81% decline**.[1]

Most striking, the email-embedded QR delivery method that surged 336% in March — flagged in our April issue as an emerging tactic — **effectively disappeared in Q2**. A delivery technique went from the fastest-growing vector in the corpus to negligible within one quarter, which is not something that happens through defensive improvement alone.

> **THE FINDING, STATED PLAINLY**
>
> Seizing the infrastructure of a single phishing-as-a-service platform reduced global QR-code phishing by 56% and CAPTCHA-gated phishing by 81% within one quarter. The concentration of capability in a small number of criminal service providers is a vulnerability in the criminal economy, and it is exploitable.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

1.3 Where the Volume Went

Our May issue forecast that displaced activity would "surface in adjacent channels rather than disappearing". Microsoft's Q2 data identifies the channel precisely, and it is not email.

Microsoft Teams became the growth surface of the quarter. Teams-based phishing volume climbed 19% from March to April, 1% from April to May, and 10% from May to June. Voice phishing over Teams grew far faster: weekly attempts rose **31% between April and May** and a further **27% between May and June**, reaching **nearly ten times the mid-2025 baseline** by the end of the quarter.[1]

Two operational details in that data are worth extracting. First, attacker display names shifted: by June, **52% used generic names** rather than obvious IT-support impersonation — an adaptation to user awareness training that keys on "IT Support" as a red flag. Second, peak activity concentrated in the **14:00–20:00 UTC window, Monday to Friday**, which is to say during business hours across European and US-Eastern working days. This is not opportunistic traffic; it is scheduled.

> **KEY TAKEAWAY**
>
> Enterprises that have spent a decade building email security maturity now face their fastest-growing social-engineering channel inside a collaboration platform that most treat as internal and trusted. Teams calls from outside the tenant deserve the scrutiny that external email already receives, and in most estates they do not get it.

1.4 Payload Composition

Credential phishing remained the overwhelming objective across the quarter at **94–96% of all payload-based attacks** — a proportion so stable it is effectively a constant.[1] What changed was the container.

PAYLOAD CONTAINER	MOVEMENT ACROSS Q2	READING
HTML attachments	**35–41%**	Stable leading container
PDF attachments	**24–31%**	Stable second
PDF share of QR delivery	**79% → 58%**	Displaced by document formats
DOC / DOCX share of QR delivery	**≈20% → 40%**	Doubled across the quarter
SVG in CAPTCHA-gated attacks	**5% → 26%**	Fivefold rise, April to June
ICS calendar files	**+277% in June**	Quadrupled in a single month
Email-embedded QR codes	**effectively zero**	Vanished after March's surge

The SVG rebound and the ICS surge are the two to watch. SVG is a text-based format that can carry script and that many gateways still treat as an image; ICS calendar files arrive with an implicit action attached and are frequently processed automatically by the client. Both are containers that carry less scrutiny than the formats they are displacing, which is the only selection criterion that has ever mattered here.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

1.5 Business Email Compromise

BEC volumes were volatile through the quarter: **9 million attacks in April** — a 121% increase over March and an outlier by any reading — followed by **3.4 million in May** (−62%) and **3.9 million in June**, which Microsoft characterises as a return to the historical baseline.[1]

The composition data is more useful than the volume. Between **87% and 92%** of initial BEC contact consisted of *generic outreach* carrying no financial request at all, with explicit financial asks appearing in only **3–8%** of messages. Invoice-payment requests fell 67% in May and a further 77% in June; payroll-update requests declined from 4% of BEC in March to 2.3% in June.

This is a meaningful behavioural shift and it degrades a widely deployed control. Detection tuned to spot financial-request language — invoice terminology, payment-detail changes, urgency markers around transfers — is looking for content that appears in fewer than one in ten initial BEC messages. The opening move is now an innocuous relationship-building exchange, and the financial request arrives only once a reply has established a thread.

1.6 Two Campaigns Worth Studying

Automated BEC at scale — 1 June

A campaign reached more than **67,000 users across in excess of 42,000 organisations in under three hours**, targeting retail and consumer goods (17%), technology (15%), and financial services (14%). Messages were generated programmatically in Python and delivered via the Amazon SES API. Critically, **neither lure contained a malicious link or attachment**.[1]

There was nothing for a content filter to detect. The messages were legitimate mail, sent from legitimate infrastructure, containing no malicious artefact — because the attack does not begin until the recipient replies. Volumetric and behavioural sender analysis is the only control that engages with this at all.

Multi-stage delivery — 14–15 June

A second campaign targeted more than **107,000 users across 19,000 organisations** using nested EML files and calendar invitations, delivered through an OAuth redirect to a collaboration-platform attachment host, terminating in a batch-file dropper that retrieved malware from a public file-sharing service.[1]

Every hop in that chain traverses a legitimate, high-reputation service. The nested EML defeats attachment inspection that does not recurse; the OAuth redirect defeats URL reputation; the file-sharing host defeats domain blocking. It is the redirect-chain technique documented in our May issue, executed at scale across four distinct trusted intermediaries.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

SECTION 2

Initial Access and Authentication Abuse

Cisco Talos published its Q2 2026 incident-response trends on 28 July, and the findings sit in productive tension with Section 1. Aggregate campaign volume fell; successful intrusion rose sharply.[2]

Incident-Response Trends — Q1 vs Q2 2026

Phishing passed half of all engagements; authentication abuse nearly doubled

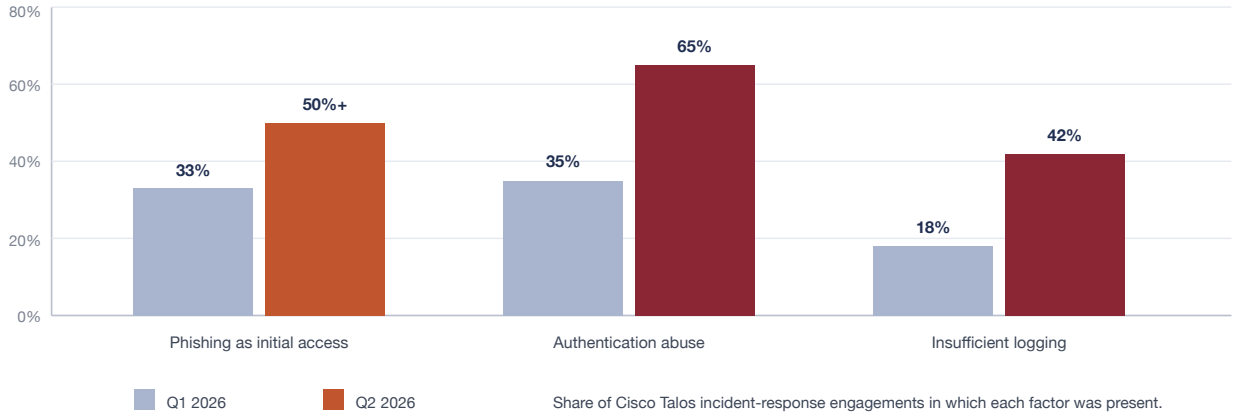

Source: Cisco Talos, IR Trends Q2 2026 (published 28 July 2026); IR Trends Q1 2026.

FIGURE 2 *Every metric moved the wrong way in the same quarter that aggregate phishing volume fell 8%. Fewer attacks, more successful intrusions — the central paradox of Q2 2026.*

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

2.1 Phishing Passes Half

Phishing was the primary initial-access method in **more than half** of Talos' Q2 engagements, up from roughly a third in Q1.[2] Exploitation of public-facing applications and drive-by compromise followed.

Delivery innovation continued in exactly the direction Section 1.4 documents: Talos observed **QR-code-embedded PDFs** used to bypass email gateways, and links hosted on trusted cloud platforms. Both appear in the payload-composition data as aggregate shifts; both appear in Talos' data as the mechanism by which specific organisations were compromised.

2.2 Authentication Abuse Nearly Doubles

The quarter's most consequential single number is that authentication abuse appeared in **65% of engagements, against 35% in Q1**.[2] Observed methods included adversary-in-the-middle proxies, session-token theft, MFA fatigue, and attacker-enrolled devices.

Read against Section 1, this is the resolution of the paradox. The Tycoon2FA takedown removed an enormous volume of *commodity* AiTM phishing — the mass-market, subscription-priced variety. It did not remove the technique, and it did not affect the operators capable of running their own infrastructure. Aggregate volume collapsed because the cheap tier disappeared. Successful intrusion rose because the capable tier was unaffected and continued improving.

> **THE SYNTHESIS OF SECTIONS 1 AND 2**
>
> Disruption removed the low end of the market, not the high end. For an organisation facing commodity campaigns, Q2 was materially safer. For one facing a competent operator, it was the worst quarter Talos has recorded for authentication abuse. Aggregate volume statistics describe the first population and say nothing about the second.

The practical conclusion has not changed since our May issue, but the evidence for it is now considerably stronger: TOTP and push-based MFA are defeated as a matter of routine, and device-bound credentials — FIDO2 and passkeys — are the only widely available control that changes the outcome. At 65% incidence, this is no longer a recommendation for high-value accounts. It is the baseline.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

2.3 The Logging Gap

Engagements hampered by insufficient logging rose from **18% in Q1 to 42% in Q2**.[2] This is the least discussed and arguably most actionable finding in the report.

A responder who cannot determine what happened cannot scope the compromise, cannot establish whether data left, and cannot confirm eviction. More than two in five engagements now run under that handicap. Some of the rise reflects the attack pattern itself — authentication abuse leaves fewer artefacts than malware execution, and identity-provider logs are frequently shorter-retention and less collected than endpoint telemetry. But the doubling in a single quarter suggests the gap is widening faster than attacker behaviour alone explains.

2.4 Sector Targeting and Tooling

Healthcare led Talos' Q2 engagements at **17%**, followed by public administration and manufacturing at **14% each**.[2] Healthcare's first-place position corroborates the ransomware leak-site data independently: BreachSense recorded healthcare as the most-targeted sector in both June and July.

Ransomware featured in **more than 20%** of engagements. Talos documented a previously undescribed variant, **Sinobi**, alongside the known Nitrogen and Warlock families. Both Sinobi and Warlock operators were observed weaponising legitimate remote management tooling — a trojanised **MeshAgent** binary and the **Zoho Assist Unattended Agent** respectively — to maintain persistent access without an active user session.

> **KEY TAKEAWAY**
>
> Remote monitoring and management tooling is now standard tradecraft for persistence. These are signed, legitimate binaries that behave exactly as intended; detection has to rest on whether the tool has any business being on that host at all. An accurate inventory of sanctioned RMM software, enforced by allowlist, is worth more here than any signature.

Taken with Section 1, Q2 2026 produced a clear operational picture: attackers get in through the human layer, defeat authentication rather than exploiting software, persist using legitimate administration tools, and operate in estates where nearly half the time the defender cannot reconstruct what happened. None of those four steps involves malware in the conventional sense.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

SECTION 3

Brand Impersonation — Q2 2026

Check Point Research published its Q2 2026 brand phishing report on 23 July. The ranking is broadly stable at the top and contains one genuinely new entry.[4]

Most-Impersonated Brands — Q2 2026

Five brands account for more than half of all brand phishing; ChatGPT enters the top ten

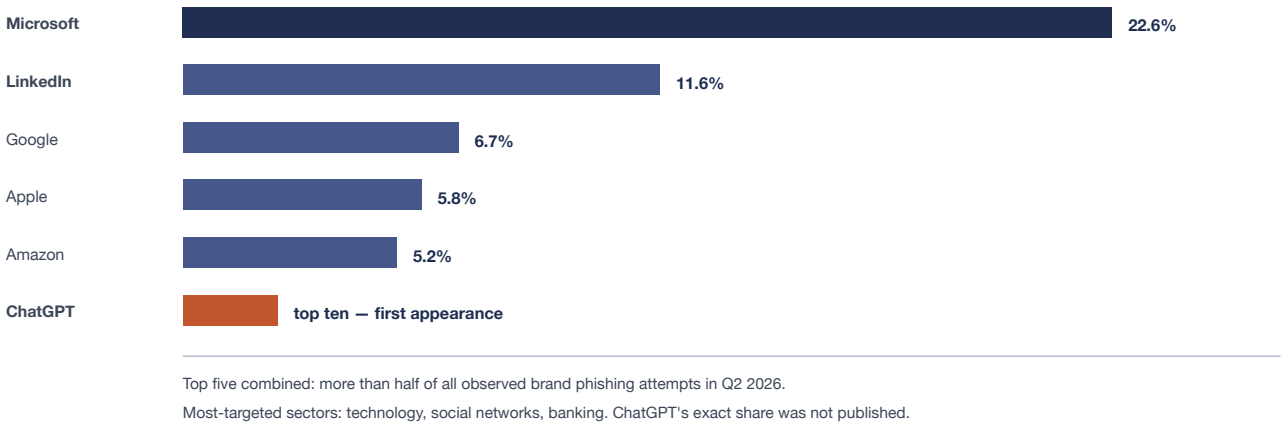

Source: Check Point Research, Q2 2026 Brand Phishing Report (published 23 July 2026). Check Point's own summary rounds the Microsoft figure to 23%.

FIGURE 3 *Microsoft's dominance is unchanged, but LinkedIn at second place is notable — a professional network, not a payment or productivity service, carrying more than one impersonation in ten.*

LinkedIn's position at **11.6%** deserves more attention than it usually receives. LinkedIn is not a payment service and holds comparatively little directly monetisable data. Its value to an attacker is *relational*: a compromised professional account yields a verified identity, a colleague graph, and a channel through which a subsequent approach carries inherent credibility. That is the raw material for the supplier-compromise pattern documented across this series.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

3.1 ChatGPT Joins the List

OpenAI's ChatGPT appeared among the ten most-impersonated brands for the first time in Q2 2026.[4] The documented campaign pattern is unremarkable in construction and significant in target selection: a fake subscription-payment-failure notice, styled with authentic OpenAI branding and subject line, leading to a payment page built to harvest full credit-card details.

Three properties make AI-service subscriptions an attractive impersonation target, and they will apply to every service in this category rather than to ChatGPT specifically:

- **Recent, individually-initiated subscriptions.** Large numbers of people signed up personally and recently, often on a corporate card, and have no established expectation of what genuine billing correspondence looks like.
- **High perceived cost of interruption.** A user who has integrated the service into daily work responds urgently to a payment failure, which is precisely the state in which scrutiny drops.
- **Weak organisational visibility.** These subscriptions are frequently outside IT's inventory, so neither the user nor the security team has a reference point for legitimate billing communications.

The prediction that follows is straightforward: as AI-service subscriptions proliferate across consumer and enterprise populations, the whole category will become a standard impersonation target. Organisations should inventory which AI services staff subscribe to and establish a billing-communication baseline for each, in the same way they already do for their major SaaS suppliers.

3.2 Sector Concentration and Technique Mix

Technology, social networks, and banking faced the highest targeting pressure in Q2.[4] Documented campaign techniques spanned fake payment-failure notices, replica online storefronts — Check Point cites imitation Michael Kors and UNIQLO sites, the latter identifiable by non-functional social links — fraudulent Apple iCloud and PayPal login pages, and fake Microsoft support pages delivering malware.

The recurring detectable weakness across the replica-storefront campaigns is instructive: **the parts of a cloned site that nobody bothers to finish**. Social media links that go nowhere, footer pages that 404, region selectors that do not switch. Automated brand protection that checks link integrity on suspected clones catches a meaningful share of these without needing to reason about visual similarity at all.

> **KEY TAKEAWAY**
>
> Brand impersonation defence has a cheap, unglamorous detection surface: cloned sites are almost never complete. Link-integrity checking on lookalike-domain candidates is a higher-yield control than most of the visual-similarity tooling sold for this purpose.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

SECTION 4

Malware and Ransomware

4.1 A 2026 High

BreachSense recorded **811 victims posted across 66 active leak sites** in July 2026, spanning 78 countries — up **15% from June's 707** and the highest monthly total of the year.[3]

Ransomware Leak-Site Postings — January to July 2026

July set a 2026 high, five weeks after the Operation Endgame takedown

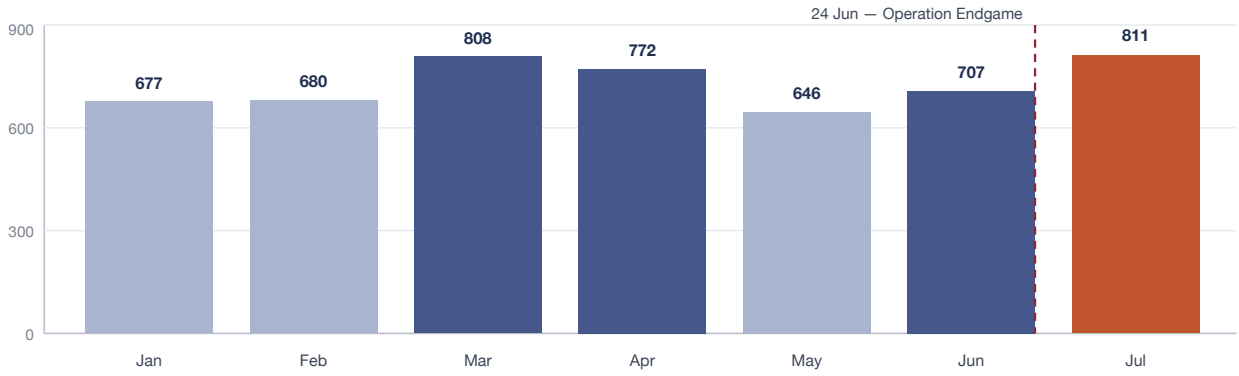

Counts are claimed victims published to leak sites, not confirmed breaches; publication lags compromise by weeks to months.

Source: BreachSense monthly ransomware reports, January to July 2026.

FIGURE 4 *The first full month after Operation Endgame produced the year's highest victim count. Given the publication lag, this does not disprove the takedown's effect — but it does rule out an immediate one.*

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

4.2 A Dead Heat at the Top

July's ranking produced an unusual outcome: **TheGentlemen and Qilin finished level at 119 victims each**. Qilin recovered strongly from June's 71 — a 68% month-on-month rise — while TheGentlemen continued its climb from 94.[3]

More instructive is what happened to June's debutant. **DeadLock collapsed from 81 victims to 22**, a 73% decline in a single month. Our June issue assessed the 81-victim debut as most likely a backlog publication — an existing operation rebranding, or affiliates migrating with in-progress engagements, publishing accumulated victims at once to establish credibility. A 73% second-month fall is consistent with exactly that: the backlog was published and the ongoing operational rate is roughly a quarter of the debut figure.

OPERATION	MAY	JUNE	JULY	READING
TheGentlemen	70	94	**119**	Sustained growth across three months
Qilin	101	71	**119**	Recovery to joint first
DeadLock	—	81	**22**	Backlog publication, then baseline

Healthcare remained the most-targeted sector and rose to **71 victims — its highest since February**.[3] This corroborates Talos' independent finding of healthcare at 17% of Q2 incident-response engagements, and the convergence of two unrelated corpora on the same sector is worth taking seriously. Healthcare combines high operational urgency, extensive third-party integration, and legacy systems that are difficult to patch — a combination that maximises both the likelihood of compromise and the pressure to pay.

4.3 What July Says About Endgame

July was the first full month after the 24 June takedown,[10] and it produced the year's highest victim count. Three readings are available, and we favour the third.

- **The takedown failed.** Unsupported: leak-site postings lag compromise by weeks to months, so July's victims were largely compromised before the action.
- **It is too early.** Partly right, but this cannot be the answer indefinitely; at some point the absence of effect becomes the finding.
- **It worked as designed, and this is not the metric.** Endgame removed access-layer infrastructure and credentials. Its effect should appear in *future* compromise rates, not in publication of past ones — and the honest position is that this question becomes answerable in September, not now.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

SECTION 5

Vulnerabilities and Major Incidents

5.1 The Largest Patch Tuesday on Record

Microsoft's June 2026 Patch Tuesday released fixes for **206 vulnerabilities**, more than **30 rated Critical** — the largest single release in the programme's history. Hornetsecurity attributes the volume in part to **AI-assisted vulnerability discovery**.[5]

If that attribution holds, it marks a structural change in patch economics that defenders have not yet adapted to. Discovery capacity has historically been the binding constraint on how many vulnerabilities get found; if AI tooling relaxes that constraint for researchers, it relaxes it symmetrically for attackers, and remediation capacity — not discovery — becomes the bottleneck that determines exposure. A 206-vulnerability month is not a manageable patch cycle for most organisations, and the triage discipline that follows from that is now a first-order security control.

Among the June release, **CVE-2026-42897** affecting on-premises Exchange Server 2016, 2019, and Subscription Edition was **actively exploited before the patch was available**. It is a cross-site scripting flaw permitting arbitrary JavaScript execution.[5] On-premises Exchange remains a high-value target precisely because the organisations still running it are disproportionately those unable to migrate — which correlates with limited patching capacity.

5.2 CISA KEV Additions — July 2026

DATE	ADDED	NOTABLE ENTRIES
7 July	**3**	Langflow (CVE-2026-55255), JoomShaper SP Page Builder, Joomlack Page Builder
10 July	**2**	iCagenda, Balbooa Forms
21–22 July	**6**	Multiple products across the catalogue

Federal remediation obligations follow from Binding Operational Directive 26-04.[6] The concentration of Joomla-ecosystem components — three separate extensions across two additions — is a reminder that CMS plugin estates remain a reliable route to web-server compromise, and that the compromised legitimate sites underpinning fake-update delivery come from somewhere.

One connection is worth drawing explicitly. Langflow appears in the 7 July KEV addition, and Hornetsecurity separately documents the **JadePuffer** ransomware operation exploiting a Langflow vulnerability (CVE-2025-3248) alongside a 2021 Alibaba Nacos flaw (CVE-2021-29441), encrypting **1,342 service configuration items** in the documented engagement.[5] A five-year-old vulnerability chained with a current one, against AI-development infrastructure, is a fair summary of where the exposure surface has moved.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

5.3 Major Incidents — July 2026

DATE	ORGANISATION	NATURE	REPORTED SCALE
16 / 27 Jul	**Coca-Cola (Fairlife)**	Ransomware; Anubis group	≈1 TB claimed exfiltrated; disclosed via SEC Form 8-K
22 Jul	**Origin Energy**	Cyber attack; ASX disclosure	Possible unauthorised access to customer data
Jul	**Suno**	Account data leak	78 million accounts

The Coca-Cola disclosure pattern is the notable one. An initial **Form 8-K** filing on 16 July was followed by confirmation on 27 July that the attack on its Fairlife dairy subsidiary had resulted in a data breach, with the Anubis group claiming encryption and roughly a terabyte of exfiltrated data.[7]

The eleven days between initial filing and confirmation illustrate how SEC materiality disclosure requirements are reshaping incident communication. Organisations now file before scoping is complete, which produces earlier public awareness — genuinely valuable — at the cost of an interval in which the only detailed narrative available comes from the attacker's leak site. Defenders consuming breach intelligence in that window are, in effect, reading the adversary's press release.

5.4 The Subsidiary Pattern

Fairlife is a Coca-Cola subsidiary, and the incident is the fifth consecutive month in which a major disclosure traced to a party other than the parent brand — contractors and OAuth integrations in April, a platform provider in May, a licensing vendor and a white-label mail platform in June, and a subsidiary here.

> **FIVE MONTHS, ONE PATTERN**
>
> Across five consecutive issues of this series, the organisation compromised has not been the organisation whose name appeared in the headline. Subsidiaries, suppliers, platform providers, and contractors are where the intrusions occur; parent brands are where the consequences land. Security programmes scoped to the parent entity are scoped to the wrong boundary, and have been all year.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

SECTION 6

vSpam Researcher Team Analysis

6.1 Corpus and Scope

24,254
INDICATORS FIRST SEEN IN JULY

1,084
THREATFOX INDICATORS — 5.1× JUNE

1,538
IP-TYPE DISTRIBUTION INDICATORS

512
NEW CLICKFIX-FAMILY LABEL

The vSpam.org collector ingested **24,254 indicators** by first-seen timestamp in July — 23,170 from URLhaus and **1,084 from ThreatFox**.[12] The headline volume continued its quarterly decline, down 14.2% from June and 19.3% from May. Underneath it, two compositional changes matter more than the total.

The first is ThreatFox. That source contributed 183 indicators in May and 214 in June; **1,084 in July is a 5.1-fold increase**. ThreatFox is oriented toward command-and-control and payload infrastructure rather than distribution URLs, so a jump of that size in that source, in a month when total volume fell, indicates a shift in what kind of infrastructure is being observed rather than how much.

INDICATOR TYPE	CLASSIFICATION	JUL	JUN	MAY
URL	Malware distribution	**15,656**	21,837	20,801
Domain	Malware distribution	**5,976**	6,236	9,055
IP	Malware distribution	**1,538**	0	0
Domain	Payload delivery	**519**	34	9
IP	Botnet command & control	**492**	155	155
Domain	Botnet command & control	**59**	11	11

The second change is visible in that table and is the more significant of the two. **IP-type malware distribution appears for the first time at 1,538 indicators**, having been entirely absent in May and June. **Domain-type payload delivery rose from 34 to 519**, a fifteenfold increase. And IP-type command-and-control, static at exactly 155 for two consecutive months, **tripled to 492**.

> **SCOPE STATEMENT**
>
> As in previous issues: community submission volume was low through this window, so nothing here derives from community reporting. Family labels are inherited from upstream feed taxonomies which mix true family names with file architectures and behavioural tags. The emergence of new indicator-type categories may partly reflect upstream feed coverage changes as well as attacker behaviour; we cannot fully separate the two and say so explicitly.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

6.2 The Technique Survives Its Operator

Our June issue[14] committed to tracking one question above all others: whether the fake-update technique would continue declining after Operation Endgame removed SocGholish, or whether it would migrate to a successor. July's data gives an answer.

ClearFake fell again, to 2,942 indicators — down 13.1% from June's 3,387 and 69.5% from May's peak of 9,632. On its own that would suggest the technique was in structural decline. But a new label entered our leading families in the same month: **`iclickfix`, at 512 indicators**, absent from May and June entirely.

LABEL	MAY	JUN	JUL	CHARACTER
clearfake	9,632	3,387	**2,942**	Fake-update family, declining
iclickfix	—	—	**512**	ClickFix-family label, new
luajit-loader	—	5,059	**—**	Single-month burst, not repeated
loader	—	—	**508**	Generic loader label, new to top 12
botnetdomain	572	269	**1,105**	Behavioural tag, quadrupled

Our reading is that **the technique is outliving the operator**. ClickFix-class social engineering — inducing the victim to execute the payload themselves — is not disappearing from our corpus; it is being relabelled as the tooling behind it changes hands. The aggregate ClickFix-family volume is falling, which is a real effect and plausibly attributable to enforcement, but the floor is not zero and a successor label is already visible.

Two competing readings deserve to be stated. The relabelling could reflect upstream taxonomy changes at URLhaus rather than genuine tooling change — a possibility we cannot exclude from our position downstream of the feed. And 512 indicators is a small number against ClearFake's 2,942; treating it as a successor requires it to grow. We will report next month whether it did.

> **JUNE QUESTION, ANSWERED — PROVISIONALLY**
>
> **Question:** does the fake-update technique decline with SocGholish, or migrate? **Answer:** both. Aggregate volume is down 69.5% from May's peak, and a new ClickFix-family label has appeared. Enforcement suppressed the technique substantially without eliminating it — which is precisely the outcome our June issue predicted for supply-chain disruption generally.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

6.3 Infrastructure Moves to the IP Layer

The compositional shift in Section 6.1 is the most consequential thing in our July corpus, and it points in a direction defenders should find uncomfortable.

Indicator Composition by Type — May to July 2026

vSpam.org collector; URL-based distribution contracts as IP-based distribution appears

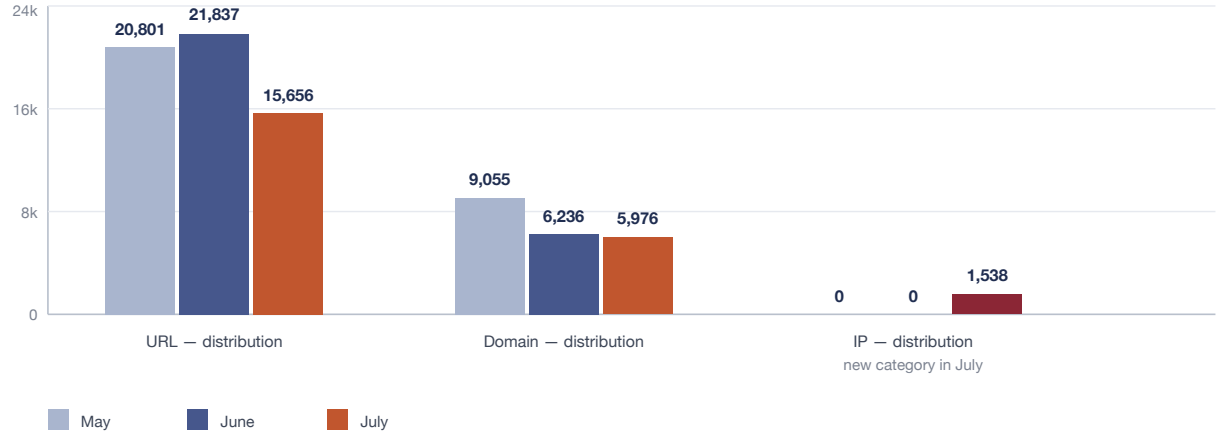

Source: vSpam.org collector corpus, indicators by type and classification, first-seen timestamp, May–July 2026.

FIGURE 5 *URL-based distribution fell 28.3% from June while an IP-based category appeared from nothing at 1,538 indicators. Domain-based distribution has declined in every month of the quarter.*

Distribution served directly from an IP address, with no domain in the chain, defeats every control that operates on names: DNS filtering, domain reputation, registrar takedown, and certificate-transparency monitoring all have nothing to act on. The trade-off for the attacker is that IPs are harder to rotate cheaply — which is why this was historically uncommon.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

That this appears in the same quarter as tightened registrar obligations and sustained domain-layer enforcement is suggestive. If domains become reliably expensive to keep — through faster takedown, registration-time verification, or registry-level detection — then serving directly from IP infrastructure becomes comparatively attractive despite its rigidity. We offer this as a hypothesis consistent with the data, not as a demonstrated causal claim, and we note that upstream feed coverage changes could produce the same appearance.

6.4 The Quarter in Our Corpus

July also produced the highest IoT and embedded-botnet reading of the quarter, at **1,351 indicators in the week commencing 27 July** — against a May–July range that had otherwise stayed between 223 and 758.

Weekly Ingestion — Fake-Update Cluster vs. IoT Botnet Baseline, Q2 2026 into July

vSpam.org collector; the two series cross in July as campaign traffic falls and infrastructure traffic rises

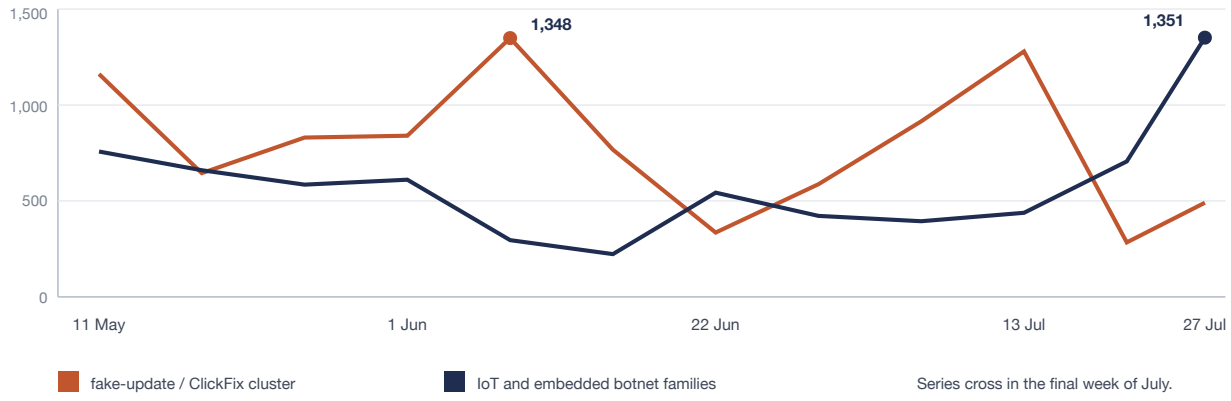

Source: vSpam.org collector corpus, weekly indicator counts by family grouping, weeks commencing 11 May – 27 July 2026.

FIGURE 6 *Twelve weeks of the two series. Campaign traffic trends down across the quarter while the infrastructure baseline holds and then spikes — the two cross for the first time in the week commencing 27 July.*

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

The crossing in the final week of July is the first time in our series that infrastructure traffic has exceeded campaign traffic in this corpus. A single week is not a trend, and IoT botnet activity is seasonal in ways we do not fully model, so we flag it as an observation to carry forward rather than as a finding.

Read across the whole quarter, our corpus describes a consistent movement. Volume declined in every month — 30,039, then 28,287, then 24,254. Within that decline, distribution moved steadily away from names and toward addresses, campaign families contracted while the infrastructure baseline held, and the domain layer shed nearly a third of its volume between May and July.

> **RESEARCHER TEAM ASSESSMENT**
>
> Our May issue[13] concluded that "the domain is becoming a disposable, bulk-acquired commodity and the URL is becoming the meaningful unit of observation". July's data extends that trajectory further than we anticipated: distribution is now appearing at the IP layer, below both. Defensive architectures anchored to DNS-layer blocking are progressively losing purchase on the infrastructure they are meant to observe — not because the blocking fails, but because there is increasingly nothing to block.

6.5 What We Will Be Watching

- Whether `iclickfix` grows past ClearFake, which would confirm succession rather than upstream relabelling.
- Whether IP-type distribution persists above 1,000 indicators per month, which would establish it as a structural shift rather than a feed-coverage artefact.
- Whether the IoT baseline spike in the final week of July continues into August, and whether the crossing observed in Figure 6 holds.
- Whether `luajit-loader` reappears — it has now been absent for a full month, which supports our June reading of it as a discrete campaign.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

SECTION 7

Forecast Scorecard

This series publishes falsifiable forecasts, and this section scores them. The Q2 retrospectives that published in July make several of the calls from our May and June issues checkable for the first time. One of them was substantially wrong.

7.1 Forecasts from the May 2026 Issue

FORECAST	OUTCOME	VERDICT
Post-takedown telemetry will show a **majority reduction** in Tycoon2FA-linked volume	92% reduction; June at ≈8% of H2 2025 baseline	**CORRECT**
Displaced activity will **surface in adjacent channels** rather than disappearing	Teams vishing reached ≈10× the mid-2025 baseline by end of Q2	**CORRECT**
Ransomware will **resume its upward trend** after May's dip	June 707 (+9.4%), July 811 (+15%, 2026 high)	**CORRECT**
The novelty-gTLD concentration will **not persist**	Cluster gone by June; `.com` resumed first place	**CORRECT**
Q2 phishing volume in the **1.0M–1.15M** range (APWG)	APWG Q2 report not yet published	**OPEN**
Telecom share falls back toward the **low-to-mid twenties**	APWG Q2 report not yet published	**OPEN**

7.2 Forecasts from the June 2026 Issue

FORECAST	OUTCOME	VERDICT
Ransomware **above 700 victims** in July and August	July 811; August pending	**ON TRACK**
Tycoon2FA step change will show **partial recovery**	No recovery — volume fell again in June, from 1.5M to 1.2M	**WRONG**
A **SocGholish successor** will emerge within one to two quarters	New ClickFix-family label at 512 indicators in July; too early to confirm	**EARLY EVIDENCE**
Sliver's share will keep growing through H2	Next Spamhaus half-year update not due until 2027	**OPEN**
Credential exposure will **outpace recovery**	78M-account leak in July against 27M recovered in June	**CORRECT**

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

7.3 The Call We Got Wrong

Our June issue forecast that Q2 telemetry would show "a majority reduction in Tycoon2FA-linked volume, with displaced activity appearing in adjacent channels rather than disappearing" — and characterised the expected shape as **a large step change with partial recovery**. The first half was right. The recovery did not happen.

Platform-linked volume fell from 1.5 million in May to **1.2 million in June**, a further 20% decline in the month when we expected the curve to begin flattening or turning.[1] Three months after the seizure, the trajectory was still downward.

Our reasoning was based on the historical pattern for criminal-service takedowns: an initial collapse followed by partial reconstitution as operators rebuild or migrate to a successor. We cited CrowdStrike's finding[11] that the platform persisted in degraded form as supporting evidence. That finding was correct — the platform does persist — but we drew the wrong inference from it, assuming persistence implied recovery. It did not.

Two explanations seem plausible, and they are not mutually exclusive:

- **Trust, not infrastructure, was the binding constraint.** A phishing-as-a-service platform whose operators have been targeted by a court-ordered seizure has a customer-acquisition problem that rebuilding servers does not solve. Criminal customers avoid services they believe are compromised or monitored, and that reputational damage does not decay on the same timescale as infrastructure.
- **Demand migrated to a different technique rather than a successor platform.** The Teams vishing surge — reaching ten times its mid-2025 baseline in the same quarter — suggests operators moved channel rather than waiting for a replacement AiTM platform. Displacement was real; it was simply lateral rather than in-place.

> **WHAT WE TAKE FROM BEING WRONG**
>
> The general lesson is that we over-weighted the infrastructure dimension of criminal service provision and under-weighted the trust dimension. Where a takedown is publicly attributed and court-ordered, the reputational damage to the service may be the more durable effect — which, if it generalises, is an argument for enforcement to maximise publicity rather than treat it as incidental.

7.4 The Central Finding of Q2 2026

Reading Sections 1 and 2 together produces the quarter's most important conclusion, and it is one that neither dataset states on its own: **aggregate phishing volume fell 8% while successful intrusion rose sharply**. Phishing passed half of all incident-response engagements; authentication abuse nearly doubled to 65%.

Enforcement removed the commodity tier of the market — the subscription-priced, mass-market tooling that generates enormous volume and comparatively few successful intrusions. It did not touch the capable tier. Volume metrics measure the tier that was removed. Incident-response metrics measure the tier that was not. An organisation reading only the first would conclude Q2 was its safest quarter in years.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

SECTION 8

Standards and Regulatory Developments

8.1 CISA

CISA made three separate additions to the Known Exploited Vulnerabilities catalogue during July 2026 — three entries on 7 July, two on 10 July, and six across 21–22 July — with federal remediation obligations following from Binding Operational Directive 26-04.[6] The Joomla-ecosystem concentration noted in Section 5.2 is the pattern worth carrying forward.

8.2 Email Authentication at Mid-Year

Continuous-measurement corpora put the mid-2026 position at **68.7% of domains worldwide with no effective DMARC protection**, **11.5%** at full protection (`p=reject` at 100% enforcement), and **19.8%** with partial coverage through quarantine or gradual rollout. Strict-policy adoption rose **2.3 percentage points over six months**.[8]

That rate of improvement — roughly 4.6 points a year — implies more than a decade to reach even half the domain population at full enforcement. Against Section 2's finding that authentication abuse appears in 65% of incident-response engagements, the more important observation is that DMARC would not have prevented most of those compromises anyway. The attacks that matter now run through genuine authenticated sessions, and no policy setting on the sending domain addresses that.

8.3 The AI-Assisted Discovery Question

Hornetsecurity's attribution of Microsoft's record 206-vulnerability June release partly to AI-assisted discovery raises a policy question this series expects to return to.[5] If discovery capacity is no longer the binding constraint on how many vulnerabilities are found, then the coordinated-disclosure model — which assumes a manageable flow into a remediation pipeline of roughly fixed capacity — comes under pressure from a direction it was not designed for.

Neither NIST nor CISA has issued guidance addressing this. The practical consequence for defenders is immediate regardless: **triage capability now matters more than patch speed**. An organisation that patches quickly but cannot rank 206 vulnerabilities by actual exposure will spend its remediation capacity in the wrong places, and the KEV catalogue is currently the best available public ranking signal.

8.4 ICANN and M3AAWG

No new milestone fell in July. The contractual regime described in our May issue and the eight proposed outcomes from M3AAWG's 67th General Meeting, reported in our June issue, remain the operative developments. Our corpus finding in Section 6.3 — distribution appearing at the IP layer — is relevant to both, since a shift below the domain layer would place infrastructure outside the reach of registrar and registry obligations entirely.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

SECTION 9

Outlook and Recommendations

9.1 Forecast (August – October 2026)

1. **Q3 phishing volume will stabilise rather than continue falling.** The Tycoon2FA effect is now largely worked through at 8% of baseline; there is little left to remove. We expect Q3 aggregate volume within 10% of Q2's 7.6 billion.
2. **Authentication abuse will exceed 65% in Q3 incident-response data.** The trend from 35% to 65% in one quarter reflects a structural shift to identity-based attack, and nothing in Q2's data suggests it has peaked.
3. **Teams and collaboration-platform vishing will keep growing.** Ten times the mid-2025 baseline, with 31% and 27% consecutive monthly rises, is a channel still early in its adoption curve.
4. **AI-service brand impersonation will expand beyond ChatGPT.** We expect at least one further AI service in a top-ten impersonation ranking by Q4 2026.
5. **Ransomware will hold above 750 victims per month through Q3.** July's 811 against a widening group count and lengthening tail suggests the market absorbed June's enforcement without material capacity loss.
6. **IP-layer distribution will persist in our corpus above 1,000 indicators monthly.** Stated so it can be scored: if this reverts to zero in August, our Section 6.3 reading was a feed artefact and we will say so.

9.2 Recommendations — Messaging Operators and ISPs

- **Score inbound by container, and reweight for SVG and ICS.** SVG rose from 5% to 26% of CAPTCHA-gated payloads across Q2 and ICS calendar files quadrupled in June. Both are formats most gateways treat as benign by type.
- **Do not rely on financial-request language for BEC detection.** Between 87% and 92% of initial BEC contact contains no financial request at all. Detection has to engage with conversational sender behaviour, not message content.
- **Add IP-direct distribution to blocking logic.** Our corpus recorded 1,538 IP-type distribution indicators in July, from a base of zero. DNS-layer controls have nothing to act on in these cases.
- **Treat aggregate volume decline with suspicion.** Q2's 8% fall coincided with a sharp rise in successful intrusion. Volume is a measure of the commodity tier only.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

9.3 Recommendations — Enterprise Defenders

— **Deploy phishing-resistant MFA as the baseline, not the exception.** At 65% incidence of authentication abuse in incident-response engagements, TOTP and push-based factors should be treated as compensating controls on their way out, not as the standard.

— **Fix the logging gap before the next incident.** Insufficient logging hampered 42% of Q2 engagements, up from 18%. Identity-provider logs — sign-in, token issuance, device registration, consent grants — are the specific telemetry that authentication-abuse investigations require, and they are the most commonly missing.

— **Extend external-communication scrutiny to collaboration platforms.** Teams calls from outside the tenant now warrant the treatment external email receives. Attackers moved to generic display names in 52% of cases by June, so awareness training keyed to "IT Support" impersonation is already behind.

— **Inventory remote management tooling and enforce an allowlist.** Trojanised MeshAgent and abused Zoho Assist deployments were documented this quarter. These are signed, legitimate binaries; the only reliable signal is whether the tool belongs on that host.

— **Build a billing-communication baseline for AI-service subscriptions.** ChatGPT's entry into the top-ten impersonated brands is the leading edge of a category effect, and these subscriptions typically sit outside IT inventory.

— **Scope security programmes to the corporate group, not the parent entity.** Five consecutive months of major incidents originating with a subsidiary, supplier, platform provider, or contractor.

9.4 Recommendations — Registrars and Registries

— **Watch for the migration below the domain layer.** If tightened registrar obligations succeed in making domains expensive, distribution moves to IP infrastructure, which sits outside registry and registrar reach entirely. Success at the domain layer should be measured against total malicious distribution, not domain-based distribution alone.

— **Maintain registration-time verification investment.** The 90% abuse reduction one registrar achieved over six months, reported in our June issue,[9] remains the strongest available evidence that this is a policy variable rather than an environmental constant.

> **WHERE WE MAY BE WRONG**
>
> The forecast most likely to fail is 9.1(6). Our IP-layer distribution finding rests on one month of data in a corpus whose upstream feed coverage we do not control, and a coverage change at the feed would produce exactly the pattern we observed. If August shows the category reverting toward zero, we will report it as an artefact and withdraw the analysis in Section 6.3.

ISSN 2026-VSPAM-007 · VOL. 2026 · ISSUE 7

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 platform telemetry	Microsoft Security, Microsoft Threat Intelligence, Barracuda, Abnormal AI, Hornetsecurity
Threat intelligence and incident response	Cisco Talos, Check Point Research, CrowdStrike, Trend Micro
Breach and ransomware tracking	BreachSense, PKWARE, UpGuard, SharkStriker, Strobes
Blocklist and reputation	Spamhaus, SURBL, URIBL, vSpam.org
Open threat-intelligence feeds	URLhaus, ThreatFox, Feodo Tracker (abuse.ch)
Remote management tooling abused this quarter	MeshAgent, Zoho Assist — cited as detection targets, not recommendations
Regulatory and coordination bodies	CISA, ICANN, M3AAWG, NIST, Europol

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.

Forecast record to date

Across the May and June issues this series made eleven scoreable forecasts. As of this issue: **six correct, one on track, one supported by early evidence, one wrong, and three still open.** The full scorecard, including our reasoning on the incorrect call, is in Section 7.

REFERENCES

References

1. Microsoft Security (23 July 2026). *Email threat landscape: Q2 2026 trends and insights.* microsoft.com/en-us/security/blog/2026/07/23/email-threat-landscape-q2-2026-trends-and-insights/
2. Cisco Talos (28 July 2026). *IR Trends Q2 2026: Phishing and weaponized remote management tools drive attack chains.* blog.talosintelligence.com/ir-trends-q2-2026/
3. BreachSense (August 2026). *July 2026 Ransomware Report: 811 Victims, 66 Groups.* breachsense.com/ransomware-reports/july-2026/
4. Check Point Research (23 July 2026). *Which Brands Are Impersonated Most? Inside the Q2 2026 Brand Phishing Report.* blog.checkpoint.com/research/
5. Hornetsecurity (July 2026). *Monthly Threat Report — July 2026.* hornetsecurity.com/en/blog/monthly-threat-report/
6. CISA (July 2026). *Known Exploited Vulnerabilities Catalog* additions of 7, 10, and 21–22 July 2026; Binding Operational Directive 26-04. cisa.gov/news-events/alerts
7. Coca-Cola Company SEC Form 8-K (16 July 2026) and subsequent confirmation (27 July 2026); see Strobes, *Top 8 Data Breaches and Exposures of July 2026,* and CM-Alliance, *Major Cyber Attacks, Data Breaches, Ransomware Attacks in July 2026.*
8. DmarcDkim.com (July 2026). *DMARC Adoption Statistics — live global data.* dmarcdkim.com/dmarc-adoption
9. Spamhaus (10 July 2026). *Botnet Threat Update, January to June 2026.* spamhaus.org/resource-hub/botnet-c-c/botnet-threat-update-january-to-june-2026/
10. Europol (24 June 2026). *Operation Endgame — coordinated action against SocGholish, Amadey and StealC.* europol.europa.eu
11. CrowdStrike (2026). *Tycoon2FA Phishing-as-a-Service Platform Persists Following Takedown.* crowdstrike.com
12. vSpam.org collector corpus, 1–31 July 2026, with comparative figures for May and June 2026. Aggregate figures reproducible via the vSpam.org public API; per-indicator data available on request to research@vspam.org.
13. vSpam.org Independent Research (June 2026). *Trends in Spam, Phishing, Spoofing, Malware and DNS Abuse — May 2026.* ISSN 2026-VSPAM-005.
14. vSpam.org Independent Research (July 2026). *Trends in Spam, Phishing, Spoofing, Malware and DNS Abuse — June 2026.* ISSN 2026-VSPAM-006.

This report is published by vSpam.org as part of the ongoing independent research series on email and DNS abuse. ISSN 2026-VSPAM-007. Volume 2026, Issue 7. Released under the Creative Commons Attribution 4.0 International License (CC BY 4.0). Citation: vSpam.org Independent Research. (August 2026). "Trends in Spam, Phishing, Spoofing, Malware, and DNS Abuse — July 2026." Volume 2026, Issue 7. ISSN 2026-VSPAM-007. Corrections, additional data, or feedback: research@vspam.org