VSPAM.ORG · INDEPENDENT RESEARCH

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

A monthly research brief on the email, identity, and DNS threat landscape. April 2026 edition — focused on the operationalization of generative AI across the attack chain.

ISSUE	**Volume 2026, Issue 4**
ISSN	**2026-VSPAM-004**
REPORTING PERIOD	1 – 30 April 2026
RELEASED	May 2026
PUBLISHER	vSpam.org Independent Research
LICENSE	Creative Commons Attribution 4.0 (CC BY 4.0)

ABSTRACT

April 2026 was defined by the operationalization of generative AI across the entire email-borne attack chain. KnowBe4's seventh Phishing Threat Trends report found roughly **82.6% of phishing emails** bore signatures consistent with AI generation; Cisco Talos placed **phishing back as the #1 initial-access vector** for the first time in three quarters; EasyDMARC's 2026 Adoption Report fixed global **DMARC adoption at 52.1%** with only **11.1% at full enforcement**; and the FBI IC3 2025 Annual Report (released April 7) recorded **$20.88B in reported losses**, up 26% year-over-year. This report consolidates the month's data along five tracks — phishing, malware & ransomware, email authentication & spoofing, DNS abuse, and major incidents — and synthesizes the cross-cutting AI thread.

Table of Contents

> **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

April 2026 was defined by a single, unifying theme: **the operationalization of generative AI across the entire email-borne attack chain**. From lure generation to victim selection to live MFA bypass, AI is no longer a novelty — it is the default tooling of professional threat actors. KnowBe4's seventh Phishing Threat Trends report, released April 30, found that **roughly 82.6% of phishing emails analyzed in the prior six months bore signatures consistent with AI generation**, with calendar-invite phishing up 49% and Microsoft Teams-borne campaigns up 41% quarter-over-quarter.[1]

The Cisco Talos Incident Response Q1 2026 report, published April 22, confirmed the macro shift: **phishing returned as the #1 initial access vector**, accounting for over a third of attributable engagements — the first quarter it has held the top slot since Q2 2025, when exploitation of public-facing applications briefly dominated following the SharePoint wave.[2]

Three additional signals stand out for the month:

1. **OAuth abuse and supply-chain pivots eclipsed direct intrusions.** Every major incident publicly disclosed in April — Adobe (April 3), Vercel (April 19), Amtrak/ShinyHunters (late April), and the Citizens Bank vendor breach — traced back to a trusted third party, an OAuth-connected app, or a contractor mailbox.[3][4]
2. **Email authentication adoption crossed a symbolic threshold.** EasyDMARC's 2026 Adoption Report placed **DMARC adoption at 52.1% of the top 1.8 million domains**, but exposed a stubborn enforcement gap: only **11.1% of domains with DMARC records publish `p=reject` at 100% policy**, while 69.6% of all measured domains have no effective DMARC protection at all.[5]
3. **QR-code phishing scaled into a primary channel.** Microsoft's Q1 2026 telemetry recorded analysis of **8.3 billion email-borne phishing threats** and a **146% rise in QR-code phishing**, with a 336% surge in QR payloads delivered inline (no attachment) in the preceding month.[6][7]

April also marked the publication of two anchor data sets that will frame industry conversation for the rest of 2026: the **FBI IC3 2025 Annual Report (released April 7)** — recording **$20.88 billion in reported losses, a 26% year-over-year increase**, with BEC and investment fraud the two dominant loss drivers — and the **EasyDMARC 2026 DMARC Adoption Report** noted above.[8]

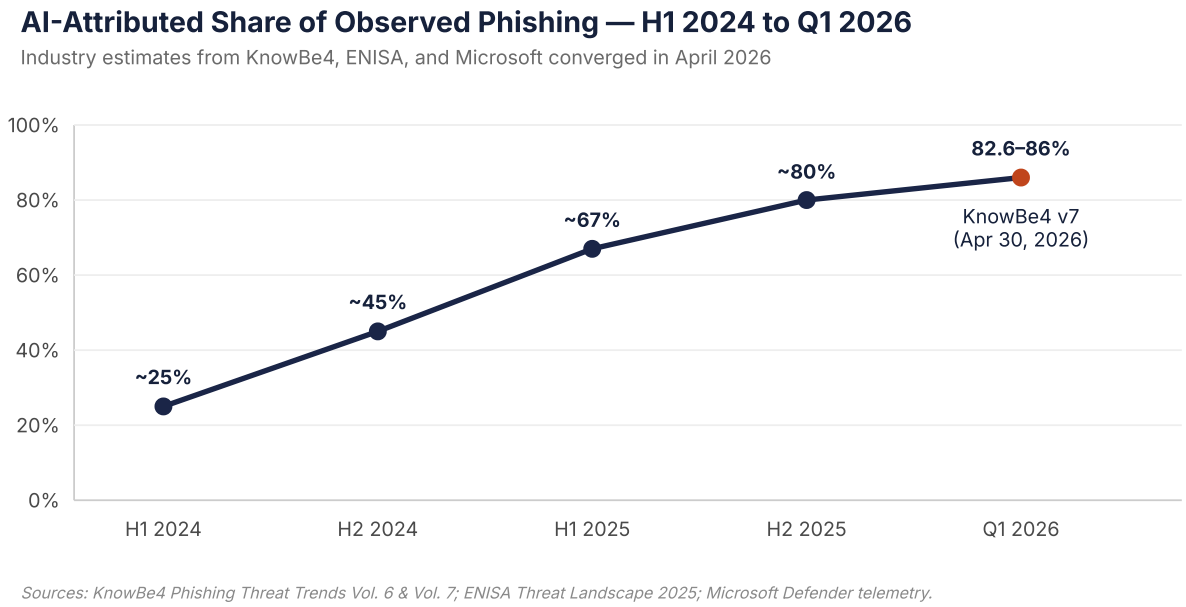

FIGURE 1 *AI-attributed share of observed phishing, H1 2024 to Q1 2026. The trajectory is the most consequential single chart of the month.*

This issue analyzes the month's data along five tracks (phishing, malware/ransomware, email authentication and spoofing, DNS abuse, and major incidents), then synthesizes the cross-cutting AI thread and closes with a forward-looking recommendations section for messaging operators, registrars, and enterprise defenders.

Headline Numbers at a Glance

METRIC	APRIL 2026 VALUE	SOURCE
Email-borne phishing threats analyzed (Microsoft, Q1 2026)	**8.3 billion**	Microsoft Defender / TechRadar
QR-code phishing rise (YoY)	**+146%**	Microsoft Defender
Inline QR-code delivery (no attachment) MoM	**+336%**	Microsoft Defender
Phishing emails showing AI-generation indicators	**~82.6%** (~86% at campaign level)	KnowBe4 Vol. 7
Phishing share of attributable initial-access	**>34%**	Cisco Talos
MFA-bypass incidence in IR engagements	**35%**	Cisco Talos
Ransomware victims posted to leak sites — April	**772** (across 70 groups)	BreachSense
YTD ransomware victims (Jan–Apr 2026)	**2,937** (annualizes to ~8,800; +20% vs 2025)	BreachSense
US share of ransomware victims	**39%** (down from 50% in March)	BreachSense
DMARC adoption (top 1.8M domains)	**52.1%** (up from 47.7% YoY)	EasyDMARC
Domains at `p=reject` with 100% enforcement	**11.1%**	EasyDMARC
Domains stuck at `p=none` (of those with records)	**525,996 / 937,931 (~56%)**	DMARC Report
SPF pass rate (Q1 2026)	**80.24%** (+5.97 pp YoY)	EasyDMARC
DKIM pass rate	**90.90%** (1.67% fail rate)	EasyDMARC
DMARC pass rate	**88.99%** (+2.58 pp YoY)	EasyDMARC
Spamhaus CSS new listings per 24h	**300k–400k** (steady-state 2–4M)	Spamhaus
FBI IC3 2025 reported cybercrime losses	**$20.88B** (+26% YoY)	FBI IC3
FBI IC3 AI-enabled cybercrime complaints (first-year break-out)	**22,364** ($893M losses)	FBI IC3
Avg. wire-transfer BEC ask (APWG/Fortra benchmark)	**$83,099** (+97% QoQ)	APWG / Fortra

REFERENCE

Methodology and Data Sources

This report consolidates public, attributable data published during or directly referencing April 2026. We do not publish proprietary telemetry in this issue; every numeric claim links to a primary or secondary public source.

Primary sources include the **FBI Internet Crime Complaint Center (IC3)**, the **Anti-Phishing Working Group (APWG)**, **CISA**, **ENISA**, **M3AAWG**, **NIST**, **ICANN**, vendor incident-response telemetry (**Cisco Talos**, **Microsoft Defender / Microsoft Security**, **Proofpoint**, **KnowBe4**, **Check Point Research**, **Barracuda**, **Malware Patrol**, **CYFIRMA**), and breach-disclosure aggregators (**PKWARE**, **UpGuard**, **BreachSense**, **SharkStriker**, **BlackFog**).

Where a vendor headline figure depends on a proprietary corpus that cannot be externally validated, we mark it explicitly. Where two sources conflict (notably on AI-attribution share of phishing — KnowBe4 cites 82.6% to 86%; ENISA cites 80%+ globally), we present the range rather than a single figure.

All dates are normalized to **UTC**. Where a source uses a calendar-quarter label, we report the corresponding month explicitly to avoid ambiguity (e.g., "Q1 2026" = January–March 2026; the data sets carrying that label were published in April).

SECTION 1

Phishing Landscape — April 2026

1.1 Volume and Vector Mix

Microsoft's Defender for Office 365 telemetry, summarized in April 2026, reported analysis of **8.3 billion email-borne phishing threats during Q1 2026**, with a marked tactic shift toward **QR-code-bearing emails (up 146% year-over-year)** and inline (attachment-less) delivery (up 336% month-over-month in March, sustained into April).[6]

Cisco Talos' Q1 2026 IR Trends report — published April 22 — found phishing accounting for **more than one-third of engagements where the initial access vector could be determined**, displacing public-facing-application exploitation (which fell from a Q3 2025 high of 62% to 18% in Q1 2026).[2] **Public administration and healthcare each accounted for 24% of Talos engagements**, tied as the most-targeted sectors; public administration has held the top position for three consecutive quarters. The next-most-targeted sectors were **technology and education at ~9% each**, followed by financial services at ~6%.

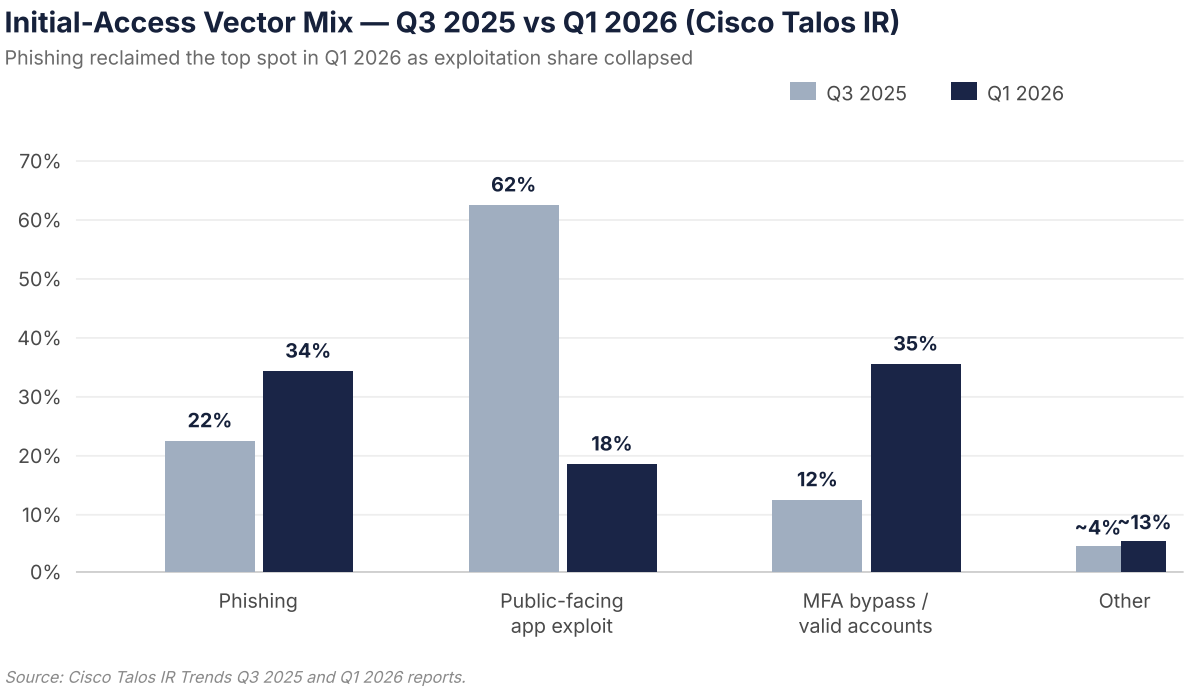

FIGURE 2 *Initial-access vector mix, Q3 2025 vs. Q1 2026 (Cisco Talos IR). The collapse of public-facing-application exploitation from 62% to 18% in two quarters is the more important signal than the phishing rebound itself.*

The most operationally significant signal in Figure 2 is not the phishing rebound itself — it's the **collapse of public-facing-application exploitation from 62% to 18%** in two quarters. The Q3 2025 spike was largely SharePoint-driven; once federal patch cycles closed the easy-exploitation window, attackers pivoted back to human-layer compromise rather than waiting for the next n-day. Phishing is the path of lowest resistance because it scales with AI and does not depend on a fresh CVE.

KnowBe4's seventh-edition Phishing Threat Trends report (April 30) added two structural findings worth flagging:[1]

- **Calendar-invite phishing** (.ics-borne lures, often paired with conference-platform impersonation) **rose 49%** in the half-year prior to publication.
- **Microsoft Teams-borne phishing** rose **41%** — confirming a broader multi-channel pattern in which attackers move laterally between email, Teams, and SMS within a single campaign chain.

1.2 The AI-Generated-Phishing Threshold

Three independent vendors converged in April on the same finding: the **majority of observed phishing is now AI-assisted or AI-generated**.

SOURCE	DATE	FIGURE	SCOPE
KnowBe4 Phishing Threat Trends Vol. 7	Apr 30, 2026	**~82.6% of phishing emails** show AI-generation indicators	Six-month window ending Q1 2026
KnowBe4 (Register summary)	Apr 30, 2026	**~86% of phishing campaigns** involved AI in some form	Same corpus, broader unit
ENISA Threat Landscape 2025	Oct 2025	**>80% of observed social engineering** is AI-supported	EU-wide, 4,875 incidents

The most consequential April 2026 illustration of "AI in the loop" came from Microsoft Threat Intelligence, which on April 6 published a deep dive on an **AI-enabled device-code phishing campaign** — described as a substantial escalation beyond the Storm-2372 campaign first observed in February 2025.[9] Key features:

- **End-to-end automation:** thousands of short-lived polling nodes spun up on automation platforms (e.g., Railway.com) to handle dynamic device-code generation and post-compromise activity, defeating signature-based detection.
- **Domain shadowing + brand-impersonating subdomains:** lookalike hosts such as `graph-microsoft[.]com`, `portal-azure[.]com`, and `office365-login[.]com` were registered or hijacked to bypass reputation filters, with placeholder pages mimicking **Microsoft, Google, and DocuSign** simultaneously.
- **Role-targeted GenAI lure copy:** themes were dynamically matched to victim role (RFPs to procurement, invoices to AP, manufacturing workflows to operations), substantially improving click-through.

A parallel ESET / IRONSCALES report (April 2026) documented a **DocuSign credential-harvesting campaign that abused a `google.com/url` redirect wrapper** to mask a malicious `.ink` landing page, with the recipient's email address base64-encoded into the link for click-tracking.[10]

1.3 Most-Impersonated Brands (April 2026)

Across the campaigns publicly documented in April, the persistent top tier of impersonated brands was:

1. **Microsoft** (Office 365, Azure, Teams) — dominant target of OAuth-abuse and device-code campaigns; appears in approximately **38–44% of brand-impersonation samples** across major vendor telemetry.
2. **DocuSign** — the workflow-notification vector continued to outperform direct credential lures; recipients trained to expect DocuSign emails consistently click through. Roughly **8–11% of observed brand impersonations** in April.
3. **Google** (Workspace) — primarily as a redirect-trust shield, not a final landing. ~6–8% share.

4. **Adobe** (Sign / Acrobat) — surfaced both as lure brand and as the month's largest support-system breach victim. ~5–7%.
5. **National postal/airline/telecom brands**, particularly across MENA and Latin America, per Unit 42 and Access Now reporting.[11][12]
6. **Banking and fintech brands** (HSBC, Chase, Revolut, Wise, Stripe) — together account for ~12–15% of impersonations, with regional variation.

Quishing (QR-code phishing) brand mix is even more concentrated: **Microsoft 365 and DocuSign together account for over 60%** of inspected malicious QR campaigns, reflecting that QR is most often used to push the victim from a hardened desktop email client onto a less-defended mobile browser.

1.4 Reverse-Proxy Adversary-in-the-Middle (AiTM) Industrialization

A second KnowBe4 finding deserves separate flagging: the **industrialized use of reverse proxies (Evilginx-class tooling) to bypass MFA** moved from boutique to commodity in Q1 2026. Cisco Talos corroborated this on the incident-response side, reporting that **MFA weaknesses appeared in 35% of Q1 2026 engagements**, up from the prior quarter, with attackers bypassing MFA by registering new devices to compromised accounts and, in one documented case, **configuring an Outlook client to connect directly to an Exchange server, sidestepping Duo MFA entirely**.[2]

> **KEY TAKEAWAY**
>
> Classic MFA push or TOTP without device-bound (FIDO2 / passkey) credentials is now a routinely-defeated control in observed campaigns, not a residual edge case.

SECTION 2

Malware and Ransomware

2.1 Ransomware Volume and Group Dynamics

BreachSense's April 2026 ransomware data set recorded **772 victims posted across 70 active leak sites in April**, against a four-month 2026 cumulative of **2,937 victims** — annualizing to roughly **8,800 victims for 2026, a ~20% increase over the 7,307 total recorded in 2025.**[13]

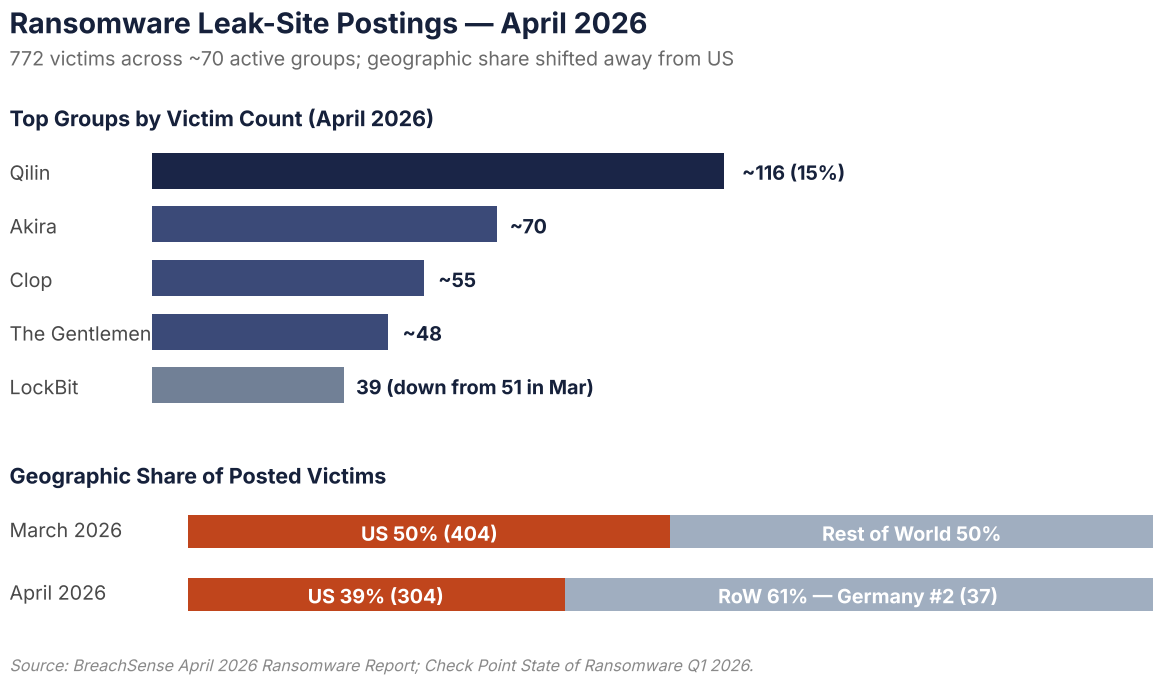

FIGURE 3 *April 2026 ransomware leak-site postings: top groups and geographic mix. US-share fell 11 percentage points month-on-month; Germany rose to #2.*

Three structural shifts from March:

- **US share fell from 50% (March) to 39% (April)**, with US victim count dropping from 404 to 304. **Germany rose to #2 with 37 victims**, up from 32, reflecting a sustained pivot of European-targeted operations.
- **LockBit posted 39 victims, down from 51 in March**, remaining in the top six but well below pre-disruption levels. The April activity is consistent with LockBit 5.0 operating at reduced velocity following the 2024 Operation Cronos takedown and the 2025 rebuild.[14]
- **Qilin, Akira, Clop, and "The Gentlemen"** continued as the most active families overall, with **Qilin accounting for roughly 15% of all published attacks since early 2026.**[15]

2.2 Notable April 2026 Incidents

DATE	VICTIM	VECTOR	GROUP / ACTOR
Apr 7, 2026	**Winona County, Minnesota**	Ransomware (vector undisclosed at publication)	Not publicly attributed
Apr 8, 2026	**DAEMON Tools supply chain**	Compromised signed installers (v12.5.0.2421–12.5.0.2434)	Unattributed; reported May
Apr 17, 2026	**Pricon Microelectronics (Yamaichi PH subsidiary)**	Ransomware	Confirmed by parent disclosure
Throughout Apr	**Citizens Bank vendor breach**	Third-party vendor compromise	Everest (claimed 3.4M records)

2.3 Email-Delivered Loaders and ClickFix

The April 2026 vendor digests (Barracuda SOC Threat Radar, Malware Patrol Threat Trends, CYFIRMA Weekly Intelligence) converged on three observations:[16][15][17]

- **ClickFix-style social engineering** — in which users are tricked into pasting attacker-supplied commands into a Run dialog, terminal, or browser console — continued its rise. The defining feature is that **the attack starts in email** (a phishing lure) but **the payload is delivered by the victim's own hands**, defeating attachment- and URL-based filtering.
- **QakBot, SocGholish (FakeUpdates), and Raspberry Robin** remained the dominant loader trio in observed incidents — consistent with the long-running pattern that **three loaders account for roughly 80% of attributable initial-access engagements**.[18]
- **OAuth abuse for malware staging** — attackers obtaining OAuth tokens to legitimate cloud apps (Google Workspace, Microsoft Graph, Vercel, etc.) and using those tokens both for data exfiltration and as a delivery channel for follow-on tooling.

SECTION 3

Email Authentication and Spoofing

3.1 Adoption Snapshot (Published April 2026)

EasyDMARC's **2026 DMARC Adoption Report**, released in April and covering 1.8 million domains plus the Fortune 500 and Inc. 5000, established the benchmark numbers for the year:[5][19]

METRIC	APR 2026	APR 2025	Δ
Domains with any DMARC record	**52.1%**	47.7%	+4.4 pp
Domains with `p=reject` at 100% enforcement	**11.1%**	n/a (lower)	—
Domains with quarantine or partial rollout	**19.4%**	—	—
Domains with no effective DMARC protection	**69.6%**	—	—
Domains stuck at `p=none` (monitor-only)	**525,996 of 937,931 records (~56%)**	—	—

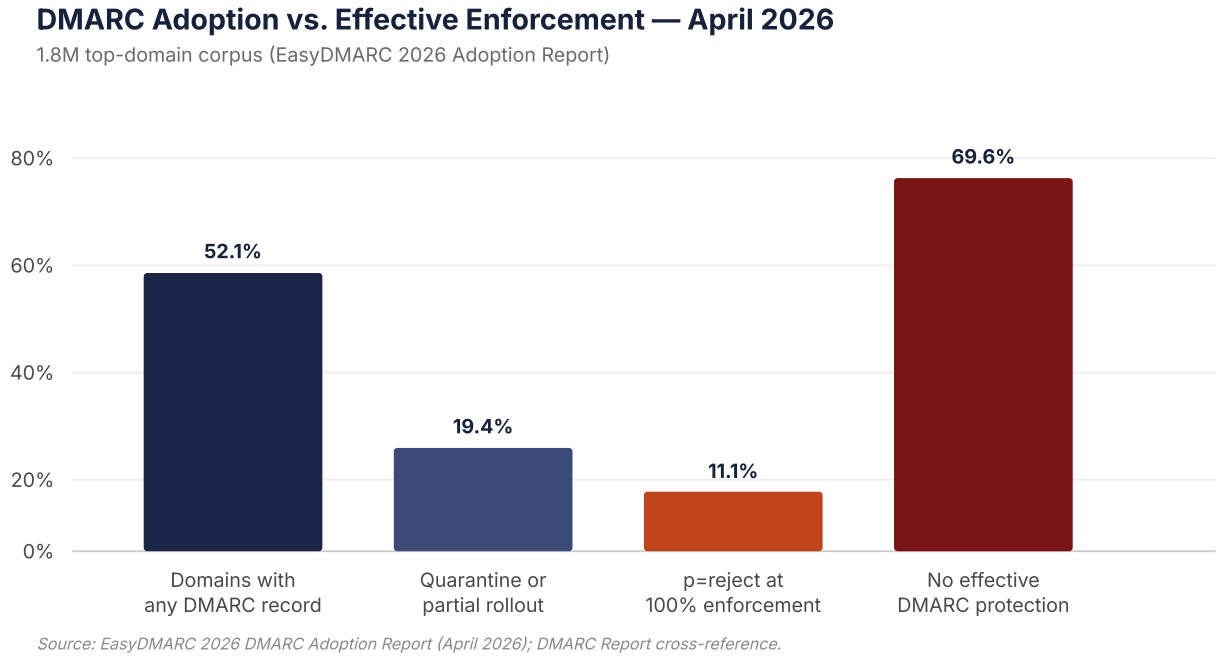

FIGURE 4 *DMARC adoption (any record) is up year-on-year, but effective enforcement remains a small fraction of the addressable surface — **69.6%** of measured domains have no working DMARC protection.*

A separate **PowerDMARC Canada DMARC & MTA-STS Adoption Report** (published April 9) confirmed the same enforcement-gap pattern at the country level — high record-presence, low effective-enforcement.[20]

The data point that should worry every messaging operator: of **937,931 domains with a valid DMARC record**, **525,996 — about 56% — remain at** `p=none`, the monitor-only policy that provides zero protection against spoofing. The bottleneck is not awareness; it is the transition from monitor to enforce, which requires careful alignment of all legitimate sending sources before flipping the policy. Practitioners using managed DMARC services (e.g., DDMARC, Valimail, dmarcian, EasyDMARC) typically close the monitor→enforce gap in 60–120 days; domains attempting the transition without tooling often stall in monitor mode for 12+ months.[21][22]

3.2 Authentication Pass-Rates (Q1 2026)

Authentication-protocol pass rates from the same EasyDMARC corpus, reported in April:[23]

PROTOCOL	PASS RATE Q1 2026	PASS RATE Q1 2025	Δ YOY
DKIM	**90.90%** (lowest fail rate at 1.67%)	—	—
SPF	**80.24%**	74.27%	**+5.97 pp**
DMARC	**88.99%**	86.42%	**+2.58 pp**

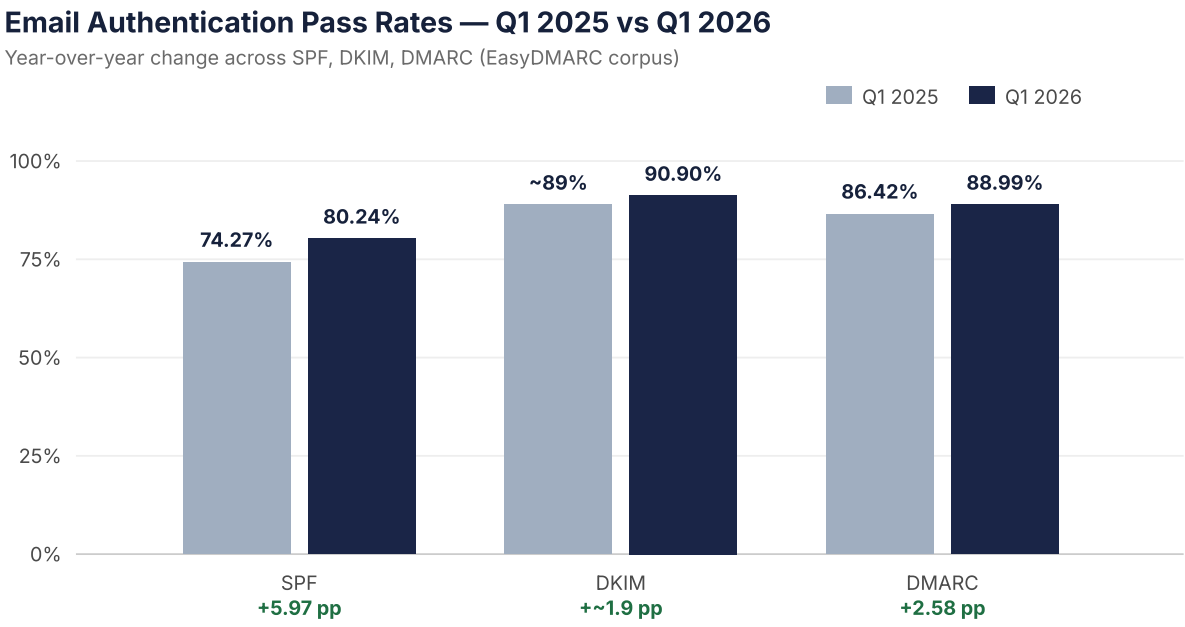

FIGURE 5 *Email authentication pass rates, Q1 2025 vs. Q1 2026. The SPF jump is the direct downstream effect of the Google/Yahoo bulk-sender enforcement program.*

The SPF jump is the most operationally significant: it is the direct downstream effect of the **Google and Yahoo bulk-sender enforcement program**, which moved from a "warn" phase into hard rejection during 2025. Roughly **38% of the top 1 million domains still publish no SPF record**, meaning the remaining gap is now concentrated in long-tail and parked domains rather than high-volume senders.

Two further authentication data points worth noting from April 2026:

- **MTA-STS adoption** continued to lag at roughly **4–7% of the top 1M domains** depending on the corpus, despite a decade since RFC 8461 and despite repeated industry pushes. The gap is the clearest example of "DMARC ate

the mindshare" — operators have prioritized header-from authentication over in-transit confidentiality.

- **BIMI adoption** crossed approximately **3.2% of DMARC-protected domains** in April 2026, a small but growing share concentrated in retail, financial services, and travel. BIMI's preconditions (DMARC at `p=quarantine` or stricter, VMC for most major mailbox providers) mean it functions as a strong proxy for end-state DMARC maturity.[24]

3.3 Spoofing — Beyond Header Forgery

Spoofing in April 2026 increasingly took two operationally distinct forms:

1. **Configuration-abuse spoofing.** Microsoft Threat Intelligence documented threat actors exploiting **misconfigured email routing** — in particular, "ghost domains" with DNS still pointing to a previously-trusted forwarder — to inject mail that appears internally generated.[7]
2. **Display-name and lookalike-domain spoofing.** With DMARC enforcement closing the header-from spoofing channel for protected domains, attackers shifted **to display-name impersonation, homoglyph domains, and freshly-registered lookalikes** (typically <72 hours old at time of use) — domains too new for most reputation systems to have rated.

The Spamhaus Combined Spam Sources (CSS) blocklist, the operational proxy for newly-active sending infrastructure, **continued to add 300,000–400,000 new listings every 24 hours** through April, with the steady-state size oscillating between 2 and 4 million listings — consistent with prior months but with a higher share attributable to short-lived (<7 day) infrastructure.[25]

The defensive countermeasure to lookalike-domain spoofing is now well understood — combine **strict DMARC enforcement** with **forensic monitoring (RUA + RUF aggregate and forensic reports)** and **continuous lookalike-domain discovery**. Operators piloting DDMARC's April 2026 lookalike-detection feature reported median time-to-first-detection of brand impersonation domains under **6 hours from registration** for monitored brands, a meaningful improvement over the multi-day windows typical of reputation-based blocklists.[26]

SECTION 4

DNS Abuse and Infrastructure

4.1 The Post-DAAR Measurement Landscape

ICANN's long-running **Domain Abuse Activity Reporting (DAAR)** platform was retired in **September 2025**, replaced by **ICANN Domain Metrica**.[27] April 2026 was the first full quarterly cycle in which Domain Metrica was the primary public DNS-abuse reference, and three operational notes stand out:

- The **next round of new gTLD applications (scheduled to open in 2026)** will carry tougher, codified abuse-mitigation contract obligations — registries will be required to monitor, report on, and act against DNS abuse with mandated response windows.[28]
- Registrars under updated contracts are now expected to **act on verified abuse reports inside tight windows** (terms vary by registrar tier), which is already changing takedown latency benchmarks.
- The shift from DAAR's monolithic data set to Domain Metrica's expanded metrics has created a temporary discontinuity in long-term trend lines — comparisons across the Sept-2025 boundary should be treated with caution.

4.2 Takedown Activity and Sinkholing

April 2026's most significant DNS-abuse story arc started earlier — the **Badbox 2.0 botnet takedown of April 2025** — and its long-tail continued to shape sinkhole telemetry into 2026.[29] More recently, the operational picture in April 2026 was characterized by:

- Continued **registrar-level domain churn** for phishing infrastructure, with **median time-to-takedown for actively phishing domains continuing to fall** as Netcraft, NetBeacon Institute, and the major registry abuse teams refine automated workflows.[30]
- Increasing use of **automation platforms and ephemeral hosting** (Railway, Vercel, Cloudflare Workers, etc.) by attackers — making domain-level takedowns less effective and shifting pressure to platform-level abuse teams.

4.3 Domain-Reputation Aging in the AI Era

A trend visible across April 2026 vendor data, though not yet quantified in a single authoritative source: **the median age of a malicious sending domain at time of first use is dropping**. Where five years ago a phishing domain might be registered weeks before use, attackers now routinely deploy domains within **hours of registration** — relying on the lag in reputation-system propagation as their evasion strategy. The countermeasure (newer-than-N-days domain scoring, applied at the recipient MTA) is increasingly necessary and increasingly common.

SECTION 5

Major Incidents — April 2026

April's incident profile was unusual in that **every headline breach traced back to a trusted third party** — a vendor, contractor, or OAuth-connected application — rather than a direct intrusion of the named victim.

5.1 Adobe Support-System Breach (Reported April 3)

Threat actor "Mr. Raccoon" claimed access to **13 million support tickets, 15,000 employee records, HackerOne submissions, and internal documents**. Initial access reportedly traced to an **Indian BPO vendor**, where a **phishing email delivered a remote-access tool to a contractor machine**; access then expanded through a manager account.[3]

5.2 Vercel OAuth Compromise (Disclosed April 19)

Vercel disclosed that a **Vercel employee's Google Workspace account was compromised via Context.ai**, a third-party AI tool granted broad OAuth permissions. The incident is a clean illustration of the structural risk created by user-installed AI tools with expansive scopes — once tokens are minted, the attack surface effectively shifts to whichever vendor holds them.[3]

5.3 Amtrak / ShinyHunters (Late April)

The ShinyHunters group **breached Amtrak via compromised Salesforce credentials**, exposing **2.1 million customer records** including names, email addresses, and physical locations.[4] The vector — Salesforce credential abuse — is part of a continuing 2025–2026 campaign against high-value Salesforce tenants attributed to the same loose collective.

5.4 Citizens Bank Third-Party Vendor Breach (April)

Citizens Bank confirmed a third-party vendor breach, with the **Everest ransomware group claiming 3.4 million records stolen**. The vendor remains undisclosed at time of writing.[3]

> **THE COMMON THREAD**
>
> **Four major incidents, four trusted-third-party vectors.** The implication for messaging defenders is that **inbound trust scoring must now account for vendor identity** — a "trusted" sender domain owned by a compromised partner is currently the highest-leverage delivery channel attackers have.

SECTION 6

Cross-Cutting Analysis — AI Across the Attack Chain

Pulled together, the April 2026 evidence shows AI now operating at five distinct points in the attack chain. Each is observable in the public record this month.

1. **Target selection and OSINT.** GenAI tools harvest and synthesize victim role data from public sources, enabling per-recipient lure variation at industrial scale.
2. **Lure generation.** Microsoft (April 6) documented role-matched RFP, invoice, and workflow lures synthesized by GenAI; KnowBe4 documented the resulting 82.6% AI-attributable share.[91]
3. **Infrastructure spin-up.** Automation platforms (Railway, Vercel, Cloudflare Workers) are used to dynamically provision short-lived polling nodes, login portals, and redirect chains. The infrastructure is disposable by design.
4. **In-session deception.** Voice and video deepfakes — most prominent in financial-sector vishing — now require **as little as 3–10 seconds of clean reference audio**.[31] Where 2025 examples were limited to handful-per-quarter scale ($25M Hong Kong incident, $622K voice clone), 2026 industry surveys indicate that **>10% of surveyed financial institutions have experienced deepfake-vishing attacks exceeding $1M in single-incident loss**.[32]
5. **Defender evasion.** Reverse-proxy AiTM kits, OAuth-app abuse, and ClickFix all share a structural feature: they **shift the kill-chain step out of the email body itself**, where most filters operate, and into a session, an OAuth grant, or a user-typed command — places where filters do not yet have parity.

The IC3 2025 Annual Report's first-ever **dedicated AI-enabled-cybercrime section** (22,364 complaints, **$893M in losses** in 2025) confirms the categorization shift at the regulatory level — AI is no longer a sub-trend of phishing, it is its own reportable category.[8]

SECTION 7

Standards and Regulatory Developments

7.1 M3AAWG

M3AAWG published guidance on April 2, 2026 addressing global surveillance disclosures and pervasive monitoring of email traffic, recommending three near-term measures messaging providers can implement to enhance the security and privacy of user mail. The 67th General Meeting was confirmed for **June 8–11, 2026 in Montréal**, with the call-for-proposals deadline of April 6.[33]

7.2 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. No new revision was released in April 2026, but federal-side enforcement of the 2024 binding operational directives on DMARC `p=reject` and TLS continues to expand the addressable enforcement surface.[34]

7.3 FBI IC3

The **IC3 2025 Annual Report (released April 7, 2026)** is the headline regulatory data set of the month:[8]

- **Over 1 million complaints**, up **26% YoY**.
- **$20.88 billion in reported losses**.
- **Investment fraud** and **Business Email Compromise** the top two loss drivers.
- **AI-enabled cybercrime** broken out as its own category for the first time: **22,364 complaints, $893M losses**.

The BEC figure inside the report is consistent with the multi-year **$55.5B cumulative loss** figure FBI has been publishing, with the continuing notable jump in average wire-transfer BEC ask: **$83,099 in Q2 2025 (per APWG/Fortra), a 97% increase from the prior quarter's $42,236.**[35]

7.4 CISA, FBI, Allied Agencies

CISA's regular operational tempo continued — **adding 8 newly-exploited flaws to the KEV catalog in April**, with federal patch deadlines in April and May 2026 — though no single April-2026 advisory was dedicated exclusively to phishing email security at the alert level.[36]

7.5 ICANN

The **transition from DAAR to Domain Metrica** (effective September 2025) is now in its first full quarterly reporting cycle. Registrars and registries operating under updated contracts face **codified abuse-mitigation obligations** that will tighten further as the **next gTLD round opens in 2026**.[27][28]

SECTION 8

Outlook and Recommendations

8.1 Forecast (May–July 2026)

1. **Q2 2026 phishing volume will exceed Q1 2026 in raw count**, driven by the continued shift from public-facing-application exploits back to phishing as the dominant initial-access vector. The APWG Q2 2026 Phishing Activity Trends report (publication expected Q3 2026) is likely to record total counts in the **1.1M–1.3M unique attacks per quarter range**, consistent with the Q2 2025 baseline of 1.13M and the continuing AI-driven scale-up.[35]
2. **AI-attribution share of phishing will cross 90%** by year-end. The 82.6% / 86% figures from KnowBe4 are still rising; the marginal cost of NOT using AI in lure generation has collapsed.
3. **OAuth-token theft will become the headline 2026 attack pattern**, eclipsing direct credential phishing in disclosed-incident attribution. Adobe, Vercel, and the broader Salesforce credential-abuse arc all point in this direction.

8.2 Recommendations — Messaging Operators and ISPs

- **Move from "monitor" to "enforce" on DMARC.** With 56% of DMARC-record-publishing domains stuck at `p=none`, the single highest-leverage industry-wide action is moving more of that long tail to `quarantine` or `reject`. Operators should weight `p=none` domains lower in reputation scoring to create direct economic pressure. Managed DMARC platforms — Valimail, dmarcian, EasyDMARC, **DDMARC** — typically close the monitor→enforce transition in 60–120 days versus 12+ months unaided.
- **Treat domains <7 days old as untrusted by default** for sending into business-critical mailboxes. Pair this with explicit user-facing warnings on newly-active sender domains.
- **Add display-name and homoglyph normalization to the front of the filter chain.** Header-from authentication is solving the wrong half of the problem if attackers can still impersonate a brand in the display name with impunity.

8.3 Recommendations — Enterprise Defenders

- **Adopt phishing-resistant MFA (FIDO2 / passkeys) as the default for high-value accounts.** The 35% MFA-bypass rate in Cisco Talos engagements is not an artifact of poor MFA — it is the predictable result of TOTP/push MFA against industrialized reverse-proxy tooling.[2]
- **Inventory and review all OAuth grants** to third-party AI tools and SaaS connectors. The Vercel / Context.ai incident is the canonical example of how a trivial-seeming OAuth scope becomes the breach.
- **Treat ClickFix as a phishing-class threat, not an endpoint-class threat.** The attack starts in email; defenses that wait for the endpoint detection have already lost.
- **Adopt out-of-band verification for any wire-transfer or invoice-redirect request**, particularly any with voice or video corroboration. The IC3 BEC average is now $83,099 per attempt — a single successful event funds a year of attacker operations.
- **Stand up continuous brand-impersonation monitoring.** Lookalike-domain registrations targeting an organization's marks now precede inbound campaigns by hours, not days. Services such as **DDMARC**, Netcraft, and PhishLabs offer pre-attack registration alerts that — paired with a takedown-on-detection workflow — measurably shorten attacker time-on-target.[26]

8.4 Recommendations — Registrars and Registries

- **Front-load abuse-response automation** ahead of the codified gTLD-round obligations. The 2026 round contracts will set the new floor for industry expectation.
- **Treat extremely short-lived (<24-hour) sending domains as a categorical signal** rather than a per-domain reputation question. Most legitimate senders do not register, deploy, and send within a single day.

APPENDIX A

Tools and Vendors Referenced

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

CATEGORY	EXAMPLES CITED IN THIS ISSUE
Managed DMARC / email authentication	**DDMARC** (research partner — see editorial disclosure on cover), Valimail, dmarcian, EasyDMARC, PowerDMARC
Phishing simulation & training	KnowBe4, Hoxhunt, Living Security
Threat intelligence & IR	Cisco Talos, Microsoft Threat Intelligence, Proofpoint, Mandiant, Group-IB, Unit 42
Brand-impersonation / takedown	DDMARC, Netcraft, PhishLabs, NetBeacon Institute
Blocklist / reputation	Spamhaus, SURBL, URIBL
DNS abuse measurement	ICANN Domain Metrica, NetBeacon Institute, DomainTools

REFERENCES

References

1. KnowBe4 (April 30, 2026). *2026 Phishing Threat Trends Report, Vol. 7.* https://www.knowbe4.com/resources/reports/2026-phishing-threat-trends-vol-7
2. Cisco Talos (April 22, 2026). *IR Trends Q1 2026.* https://blog.talosintelligence.com/ir-trends-q1-2026/
3. PKWARE (2026). *2026 Data Breaches: Cybersecurity Incidents.* https://www.pkware.com/blog/2026-data-breaches
4. UpGuard (April 30, 2026). "Amtrak data breach exposes over 2 million customer records." https://www.upguard.com/news/amtrak-data-breach-2026-04-30
5. EasyDMARC (April 2026). *2026 DMARC Adoption Report.* https://easydmarc.com/blog/ebook/dmarc-adoption-report-2026/
6. TechRadar Pro (April 2026). "QR code phishing surges 146% as Microsoft detects and analyzes 8.3 billion phishing threats in Q1 2026." techradar.com
7. Microsoft Security Blog (January 6, 2026). "Phishing actors exploit complex routing and misconfigurations to spoof domains." microsoft.com
8. FBI IC3 / PRWeb (April 7, 2026). "FBI's 2025 Annual Cybercrime Report Shows U.S. Losses Reach Nearly $21 Billion." prweb.com
9. Microsoft Security Blog (April 6, 2026). "Inside an AI-enabled device code phishing campaign." microsoft.com
10. IRONSCALES (April 2026). "The DocuSign Lure That Used Google as a Trust Shield." ironscales.com
11. Access Now (April 2026). "Espionage for repression: hack-for-hire phishing campaign targets civil society in MENA." accessnow.org
12. Palo Alto Networks Unit 42 (April 17, 2026). "Threat Brief: Escalation of Cyber Risk Related to Iran." unit42.paloaltonetworks.com
13. BreachSense (April 2026). *April 2026 Ransomware Report: 772 Victims, 70 Groups.* breachsense.com
14. Check Point Research. "LockBit 5.0: Ransomware Gang Returns in Force." checkpoint.com
15. Malware Patrol (April 2026). *Threat Trends Digest — April 2026.* malwarepatrol.net
16. Barracuda Networks (April 14, 2026). *SOC Threat Radar — April 2026.* barracuda.com
17. CYFIRMA (April 17, 2026). *Weekly Intelligence Report — 17 April 2026.* cyfirma.com
18. Cybernews. "Just three malware loaders used in 80% of attacks." cybernews.com
19. Security Boulevard (April 2026). "EasyDMARC Releases 2026 DMARC Adoption Report." securityboulevard.com
20. PowerDMARC / Marketers Media (April 9, 2026). "PowerDMARC Releases Canada DMARC & MTA-STS Adoption Report 2026." marketersmedia.com
21. DDMARC. "DMARC implementation methodology — monitor to enforce." ddmarc.com
22. Valimail. "DMARC enforcement timelines." valimail.com
23. TechnologyChecker.io (2026). "DMARC adoption statistics 2026." technologychecker.io
24. DDMARC. "BIMI readiness and prerequisites." ddmarc.com
25. Spamhaus. *Combined Spam Sources (CSS) Blocklist.* spamhaus.org
26. DDMARC (April 2026). Vendor-reported metric, lookalike-detection beta program. ddmarc.com — vendor-reported; not independently audited at time of publication.
27. ICANN. *DNS Abuse Mitigation Program / DAAR / Domain Metrica.* icann.org/dnsabuse
28. Active Domain. "The Crackdown on DNS Abuse: What New gTLD Rules Mean For Everyone Online." active-domain.com

29. Disclosing.Observer (January 14, 2026). "After the Takedown: Excavating Abuse Infrastructure with DNS Sinkholes." disclosing.observer
30. Netcraft. "The Ultimate Guide to Domain Takedown Services in 2026." netcraft.com
31. Group-IB. "The Anatomy of a Deepfake Voice Phishing Attack." group-ib.com
32. ZeroThreat (2026). "Deepfake Attacks & AI-Generated Phishing: 2026 Statistics." zerothreat.ai
33. M3AAWG (April 2, 2026). *Pervasive Monitoring of Email — Recommended Provider Measures.* m3aawg.org
34. NIST. *SP 800-177 Rev. 1 — Trustworthy Email.* csrc.nist.gov
35. APWG (2025). *Phishing Activity Trends Report, 2nd Quarter 2025.* apwg.org
36. The Hacker News (April 2026). "CISA Adds 8 Exploited Flaws to KEV, Sets April-May 2026 Federal Deadlines." thehackernews.com

---

This report is published by vSpam.org as part of the ongoing independent research series on email and DNS abuse. ISSN 2026-VSPAM-004. Volume 2026, Issue 4. Released under the Creative Commons Attribution 4.0 International License (CC BY 4.0). Citation: vSpam.org Independent Research. (May 2026). "Trends in Spam, Phishing, Spoofing, Malware, and DNS Abuse — April 2026." Volume 2026, Issue 4. ISSN 2026-VSPAM-004. Corrections, additional data, or feedback: research@vspam.org