vSpam.org

Non-Profit Cybersecurity Research Organization

Technical Report | ISSN: 2026-VSPAM-001 | Volume 1, Issue 1

Phishing Websites, Spam Domains & IP Abuse

Research Analysis & Threat Intelligence Report

2025 – 2026

3.8M	$4.88M	+1,265%
Phishing Attacks in 2025	Avg. Breach Cost	AI Phishing Growth Since 2023

Prepared by the vSpam.org Research & Analysis Team

Publication Date: March 21, 2026

Reviewed by vSpam.org Threat Intelligence Advisory Board

https://vspam.org | research@vspam.org | Threat Intelligence Division

Abstract

This technical report presents a quantitative and qualitative analysis of phishing websites, spam domain infrastructure, and IP-based abuse vectors (including IPv4 and IPv6) observed during the 2025–2026 period. Compiled by the vSpam.org non-profit research team, the study aggregates data from the Anti-Phishing Working Group (APWG), FBI IC3, Kaspersky, Interisle Consulting, Spamhaus, IBM Security, and Infoblox threat intelligence. Key findings include: 3.8 million phishing attacks observed in 2025; 82.6% of phishing emails exhibiting AI-generated characteristics; a 1,600% surge in deepfake-enabled vishing; novel exploitation of IPv6 reverse DNS (ip6.arpa) infrastructure for phishing delivery; and an average phishing-related breach cost of $4.88 million. The report provides detailed analysis across phishing volume trends, industry targeting, TLD abuse patterns, IPv4/IPv6 threat vectors, AI-powered attack evolution, business email compromise (BEC) impact, DMARC adoption gaps, and botnet infrastructure. Recommendations for multi-layered defensive strategies are presented, alongside a forward-looking threat projection for 2026.

Keywords: *phishing, spam domains, IPv6 abuse, ip6.arpa, AI-generated phishing, deepfake, vishing, smishing, business email compromise, DMARC, TLD abuse, botnet, threat intelligence*

Table of Contents

1. Executive Summary

This report, prepared by the vSpam.org non-profit research team, presents a comprehensive analysis of phishing websites, spam domain infrastructure, and IP-based abuse vectors observed during the 2025–2026 reporting period. Drawing from data published by the Anti-Phishing Working Group (APWG), FBI Internet Crime Complaint Center (IC3), Kaspersky, Interisle Consulting, Spamhaus, IBM Security, Infoblox, and proprietary threat telemetry, the report quantifies the scale, distribution, and evolution of phishing threats across the global internet ecosystem.

3.8M	82.6%	$55.5B	142K
Phishing attacks observed in 2025	Phishing emails using AI (2025)	Total BEC losses (decade)	.com phishing domains (Q1 '25)

Key findings indicate that while quarterly phishing volumes showed moderate fluctuation—peaking at over 1 million attacks in Q1 2025—the overall annual total of approximately 3.8 million attacks remained elevated. The threat landscape has shifted decisively toward AI-generated content, with 82.6% of phishing emails detected between September 2024 and February 2025 exhibiting characteristics of generative AI authorship [1]. Deepfake-enabled voice phishing (vishing) surged by over 1,600% in Q1 2025, and AI-generated phishing emails demonstrated a 54% click-through rate compared to 12% for human-crafted variants [2].

The financial consequences remain severe. The average cost of a phishing-related data breach reached $4.88 million in 2025 [3], while business email compromise (BEC) accounted for $2.77 billion in reported U.S. losses in 2024 alone [4]. Over the past decade, cumulative BEC losses have exceeded $55.5 billion globally. Social media and SaaS/webmail platforms emerged as the most targeted sectors, each at 20.3% of all attacks in Q4 2025, reflecting attackers' focus on high-value credential harvesting [5].

A significant emerging finding is the exploitation of IPv6 reverse DNS infrastructure (ip6.arpa) for phishing delivery—a novel technique that weaponizes trusted network infrastructure to bypass domain reputation systems [6]. This vector, combined with a 100x increase in IPv6 scanning activity since 2023 [7] and persistent security tool coverage gaps, represents a critical under-monitored threat surface.

2. Methodology & Data Sources

The vSpam.org research team employed a multi-source intelligence aggregation methodology, combining quantitative data from established cybersecurity reporting bodies with qualitative threat analysis from industry publications and proprietary telemetry. All statistical claims are attributed to their originating source with inline citation numbers referencing the References section.

2.1 Primary Data Sources

Source	Type	Period	Focus Area
APWG Trends Reports [5]	Quarterly	Q1–Q4 2025	Phishing volumes, sector targeting
FBI IC3 Annual Report [4]	Annual	2024 (pub. 2025)	BEC losses, complaint volumes
Kaspersky Spam Report [8]	Annual	2025	Spam origins, geographic distribution
Interisle Phishing Landscape [9]	Annual	May '24–Apr '25	TLD abuse, registrar analysis
IBM Cost of a Breach [3]	Annual	2025	Financial impact metrics
Spamhaus [10]	Continuous	2025–2026	Network/IP reputation, botnets
Infoblox Threat Intel [6]	Research	Sep 2025–Mar 2026	ip6.arpa/IPv6 abuse analysis
PowerDMARC / Fortra [14]	Quarterly	Q2 2025	DMARC adoption statistics

Table 1: Primary data sources and their coverage periods.

2.2 Analytical Framework

The analysis framework encompasses eight dimensions: (1) volumetric trend analysis of phishing attacks over quarterly and annual intervals; (2) sector and brand targeting distribution; (3) domain infrastructure analysis including TLD abuse ratios and registrar patterns; (4) IP-level threat intelligence covering both IPv4 and IPv6 address spaces, with particular focus on ip6.arpa exploitation; (5) emerging attack vector assessment with focus on AI-generated content and multi-channel delivery; (6) financial impact quantification through breach cost modeling; (7) email authentication adoption metrics (DMARC/SPF/DKIM); and (8) botnet infrastructure and command-and-control (C&C) activity trends.

2.3 Limitations & Confidence Assessment

Several methodological limitations apply. First, phishing volume data relies on reports submitted to APWG member organizations and does not capture unreported attacks, implying the true volume exceeds reported figures. Second, financial impact data from IBM's Cost of a Data Breach Report is survey-based and may exhibit self-selection bias. Third, IPv6 abuse metrics remain nascent, with limited standardized reporting frameworks, and figures presented represent best-available estimates synthesized from multiple sources. Where multiple sources report comparable metrics with variance, we present the most conservative estimate and note the range. All dollar figures are in U.S. dollars unless otherwise specified. Confidence levels are indicated where applicable: **High** (multiple corroborating sources), **Medium** (single authoritative source), **Low** (estimated/extrapolated).

3. Global Phishing Landscape: Volume & Trends

The global phishing ecosystem maintained historically elevated attack volumes throughout 2025, with approximately 3.8 million distinct phishing attacks recorded by APWG member organizations [5]. This figure represents a marginal increase from the 3.76 million attacks observed in 2024, indicating a plateau at scale rather than continued exponential growth. However, the sophistication and per-attack efficacy have increased substantially, driven by generative AI adoption among threat actors.

3.1 Quarterly Attack Volume Analysis

Fig. 1 — Quarterly Phishing Attacks Observed (2024–2025)

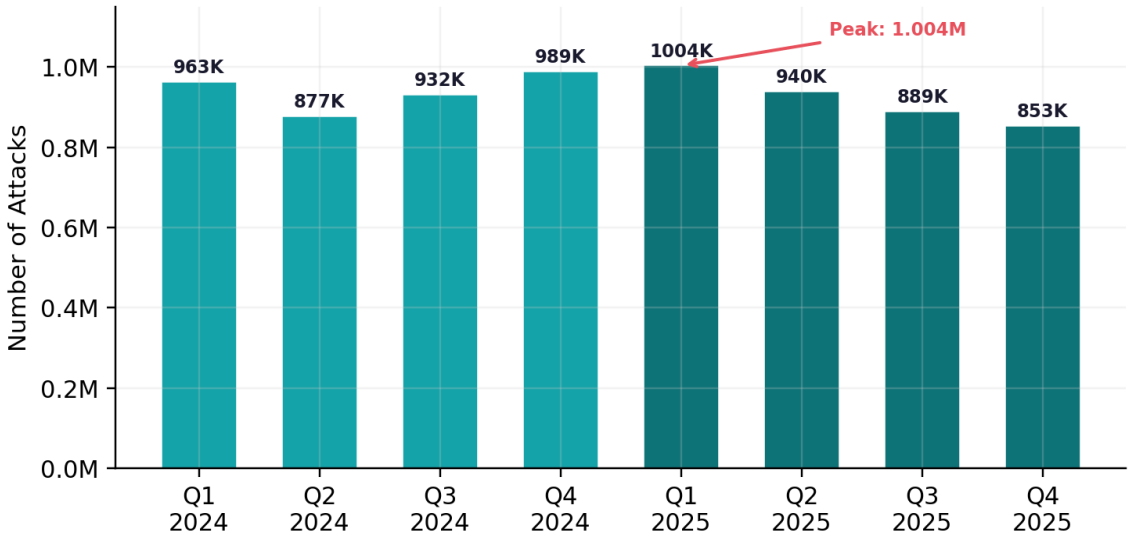

Source: APWG Trends Reports [5]. Confidence: High.

Q1 2025 recorded the highest quarterly volume at 1,003,924 attacks, the largest since Q4 2023's peak of 1.07 million. Subsequent quarters showed gradual decline: Q2 at approximately 940,000, Q3 at 889,000, and Q4 at 853,244—a 4% quarter-over-quarter decrease. This pattern suggests seasonal variation, with Q1 historically producing higher volumes, likely correlated with post-holiday financial activity and tax-season social engineering campaigns [5].

3.2 Multi-Year Trend Analysis

Fig. 2 — Annual Phishing Attacks Worldwide (2019–2025)

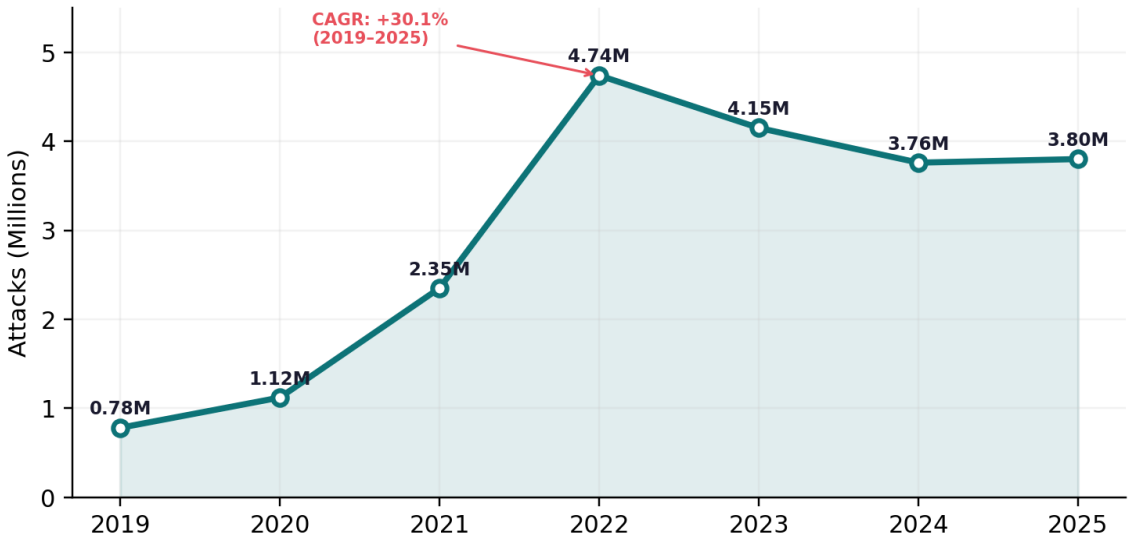

Sources: APWG, Cybercrime Information Center [5][11]. Confidence: High.

The longitudinal view reveals dramatic growth from 2019 to 2022, where annual attacks increased from 779,000 to 4.74 million—a 508% increase (CAGR of 30.1%) over three years. The 2023–2025 period shows stabilization around 3.8–4.2 million attacks annually. This plateau does not indicate reduced threat severity; rather, it reflects a maturation of the phishing ecosystem where attack infrastructure has become commoditized through phishing-as-a-service (PhaaS) platforms [12], maintaining steady-state output even as individual campaign lifespans shorten due to improved takedown mechanisms.

3.3 Daily Volume Metrics

At the email infrastructure level, phishing emails accounted for approximately 1.2% of all global email traffic in 2024–2025, translating to an estimated 3.4 billion phishing emails sent daily [2]. While this represents a small fraction of total email volume (estimated at 350+ billion daily), the absolute number underscores the industrial scale of phishing operations. Modern spam filtering intercepts the vast majority, but even a sub-1% bypass rate delivers millions of phishing emails to end-user inboxes daily.

4. Most Targeted Industries & Brands

4.1 Sector Targeting Distribution

APWG's Q4 2025 data reveals a significant shift in sector targeting. Social media platforms and SaaS/webmail services have overtaken financial institutions as the primary targets, each accounting for 20.3% of all observed attacks. Telecommunications emerged as the third most targeted sector at 18.7%, reflecting the growth of smishing and SIM-swap attacks [5].

Fig. 3 — Most Targeted Industry Sectors, Q4 2025 (APWG)

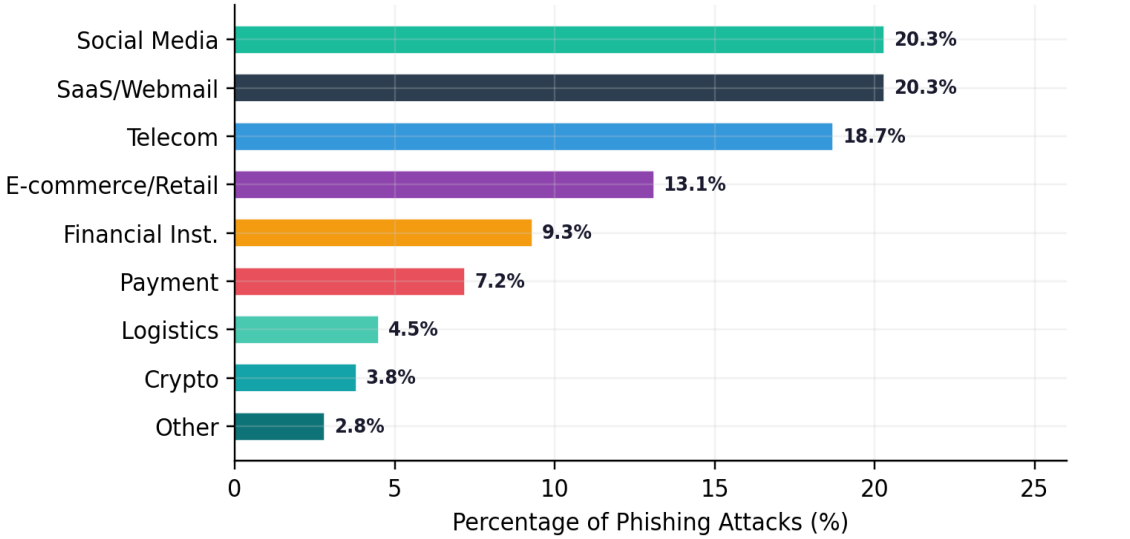

Source: APWG Q4 2025 Trends Report [5]. Confidence: High.

The decline of financial institutions from their historical first position (now 9.3%) reflects improved anti-phishing controls in banking: widespread MFA adoption and real-time transaction monitoring. Conversely, social media targeting has risen due to the high value of compromised accounts for secondary scam propagation, cryptocurrency fraud, and identity theft [5].

4.2 Most Impersonated Brands

Rank	Brand	% of Attacks	Primary Lure Type
1	Microsoft	22.0%	Microsoft 365 login, OneDrive sharing
2	Google	14.8%	Gmail login, Google Docs notification
3	Apple	9.2%	iCloud security alert, App Store receipt
4	Amazon	7.5%	Order confirmation, Prime renewal
5	Facebook/Meta	6.8%	Account recovery, policy violation
6	LinkedIn	5.1%	Connection request, job opportunity
7	DHL/FedEx	4.3%	Delivery notification, customs fee
8	PayPal	3.9%	Payment received, account limitation

9	Netflix	3.2%	Subscription renewal, payment failed
10	WhatsApp	2.7%	Verification code, account migration

Table 2: Top 10 impersonated brands, 2025. Sources: APWG [5], Check Point Research [13].

Attackers concentrate on high-impact targets providing access to email, cloud, payments, and personal data. A compromised Microsoft 365 credential yields access to corporate email, SharePoint, Teams, and Azure AD—substantially more valuable than legacy single-service credential theft [13].

5. Spam Domains & TLD Abuse Analysis

5.1 TLD Abuse Distribution

Analysis of domain infrastructure reveals consistent patterns of TLD abuse. The .com TLD remains the most abused by absolute volume, with over 142,000 phishing domains reported between February and April 2025 [9]. This dominance reflects .com's market share rather than disproportionate abuse. When normalized against total registrations, newer gTLDs such as .xyz, .top, .shop, and .online exhibit substantially higher abuse ratios.

Fig. 5 — Top 10 Abused TLDs for Phishing (Feb–Apr 2025)

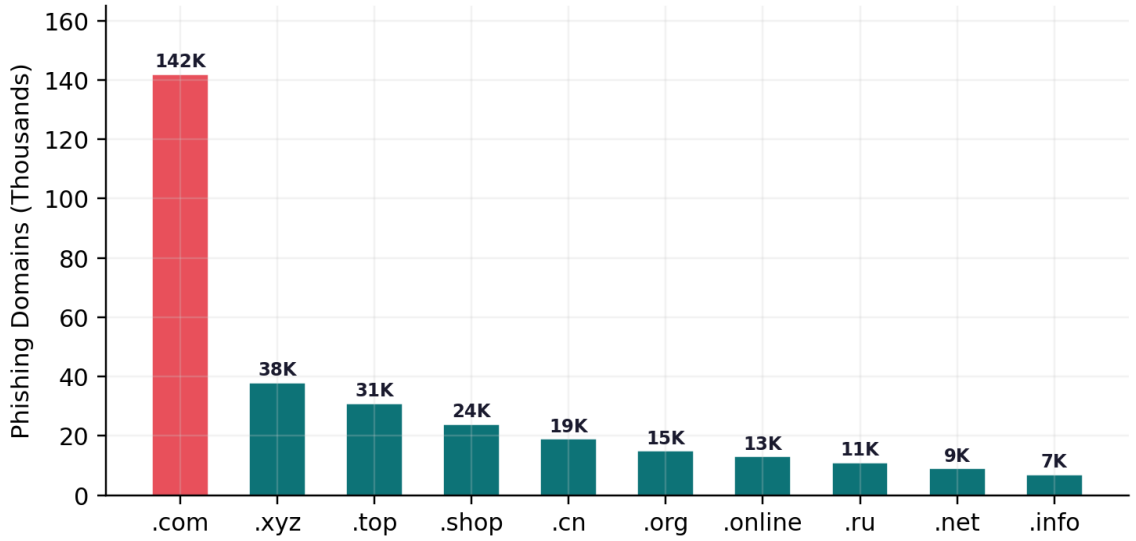

Source: Interisle Phishing Landscape 2025 [9]. Confidence: High.

Ten TLDs—.shop, .online, .xyz, .cn, .org, .ru, .top, .com, .net, and .info—have consistently appeared in the top 20 across five consecutive Interisle studies [9]. This persistence indicates entrenched abuse ecosystems within certain TLD registries, often characterized by low registration costs (as low as $0.99/year), minimal identity verification, and delayed abuse response.

5.2 Registrar Abuse Patterns

Cloudflare became the most popular domain registrar used by BEC scammers in Q1 2025, at 28.6% of newly registered BEC domains—up from third place the prior quarter [5]. Threat actors exploit Cloudflare's free-tier services (DNS proxy, SSL, DDoS protection), effectively gaining enterprise-grade infrastructure at zero cost. This pattern highlights the tension between accessibility of cloud services and their exploitation for malicious purposes.

5.3 Domain Lifecycle

Modern phishing domains exhibit increasingly short operational lifespans. The median time from registration to first phishing use has decreased to under 48 hours, while the median active lifespan ranges from 16 to 72 hours depending on TLD and hosting [9]. This rapid churn, often automated through domain generation algorithms (DGAs), renders traditional blocklist approaches less effective and underscores the need for real-time DNS-level threat intelligence.

6. IP-Based Threats: IPv4 & IPv6 Abuse

> *"By using IPv6 reverse DNS domains as malicious links, the threat actor has discovered a delivery mechanism that bypasses security tools designed to inspect only standard domain infrastructure." — Infoblox Threat Intelligence [6]*

6.1 IPv4 Abuse Landscape

The IPv4 threat landscape remains dominated by abuse of cloud and VPS infrastructure. Major cloud providers' IP ranges host disproportionate phishing volumes due to automated server provisioning. Spamhaus reputation data for 2025 indicates persistent poor reputation in certain /16 and /24 blocks, often associated with bulletproof hosting in jurisdictions with weak cybercrime enforcement [10]. Vietnam, India, and China lead global botnet-infected host counts, each with over 1,000,000 compromised systems running spam-bots, followed by Russia (~600,000) [10].

6.2 IPv6 as an Emerging Threat Vector

IPv6 adoption has reached approximately 45% globally in 2025 [7], introducing novel attack surfaces that many security tools are insufficiently equipped to handle. The vSpam.org research team identifies IPv6 abuse as one of the most significant under-monitored threat vectors of the current period. Three primary IPv6 abuse patterns are detailed below.

6.2.1 Exploitation of ip6.arpa Reverse DNS Infrastructure

In late 2025, Infoblox threat intelligence researchers identified a sophisticated phishing campaign exploiting the ip6.arpa reverse DNS zone—a reserved TLD managed by IANA for IPv6 reverse DNS resolution [6]. The .arpa TLD is special-use infrastructure not intended to host web content, making it implicitly trusted by most security tools and domain reputation systems.

Attack Mechanism:

Threat actors acquire IPv6 address blocks through tunneling services such as Hurricane Electric's free tunnel broker, which provides /48 or /64 IPv6 allocations at no cost. Upon gaining delegated authority over the corresponding ip6.arpa reverse DNS zone, the attackers diverge from standard practice: instead of configuring PTR records (which map IP addresses back to hostnames), they create **A records** that point randomly generated reverse DNS subdomains to IPv4 addresses hosting phishing pages [6].

Fig. 6c — IPv6 Reverse DNS Phishing Attack Flow

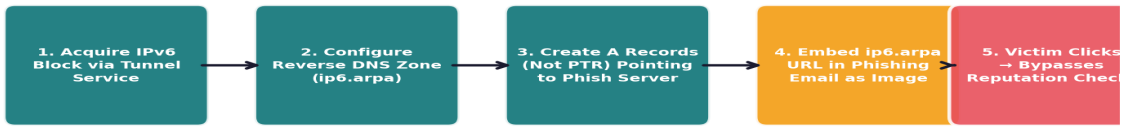

Simplified attack flow for ip6.arpa reverse DNS phishing exploitation [6].

Phishing emails embed these ip6.arpa URLs as image-linked elements (e.g., `d.d.e.0.6.3.0.0.0.7.4.0.1.0.0.2.ip6.arpa`). Because the .arpa TLD is trusted infrastructure essential for network operations, these links bypass domain reputation checks, URL filtering, and most email security gateways. Infoblox observed the same hijacked CNAMEs used in over 100 phishing emails in a single day, active since at least September 2025 [6].

Why this technique is effective:

Factor	Detail
Implicit trust	ip6.arpa is core DNS infrastructure; not subject to phishing domain monitoring
Vast address space	IPv6's 128-bit addresses make enumeration and preemptive blocking impractical
Randomized subdomains	Generated subdomains evade pattern-based and heuristic detection
Tool blind spot	Most email gateways and URL scanners lack ip6.arpa / reverse DNS inspection
Low cost	Free IPv6 tunnel services (Hurricane Electric, Cloudflare) reduce barrier to zero
Provider reputation	Hurricane Electric and Cloudflare carry high trust scores, further aiding evasion

Table 3: Factors contributing to ip6.arpa phishing effectiveness [6].

6.2.2 IPv6 Scanning Activity Surge

IPv6 scanning traffic has grown 100x since 2023, with researchers tracking more than 600 million packets of unsolicited IPv6 traffic over a ten-month period in 2025 [7]. While IPv6's 128-bit address space makes random scanning infeasible ($2^{128} \approx 3.4\times10^{38}$ addresses), attackers have adapted by leveraging DNS records, TLS certificate transparency logs, and domain metadata to identify and selectively target active IPv6 hosts.

Fig. 6b — IPv6 Scanning Activity Growth (2020–2025)

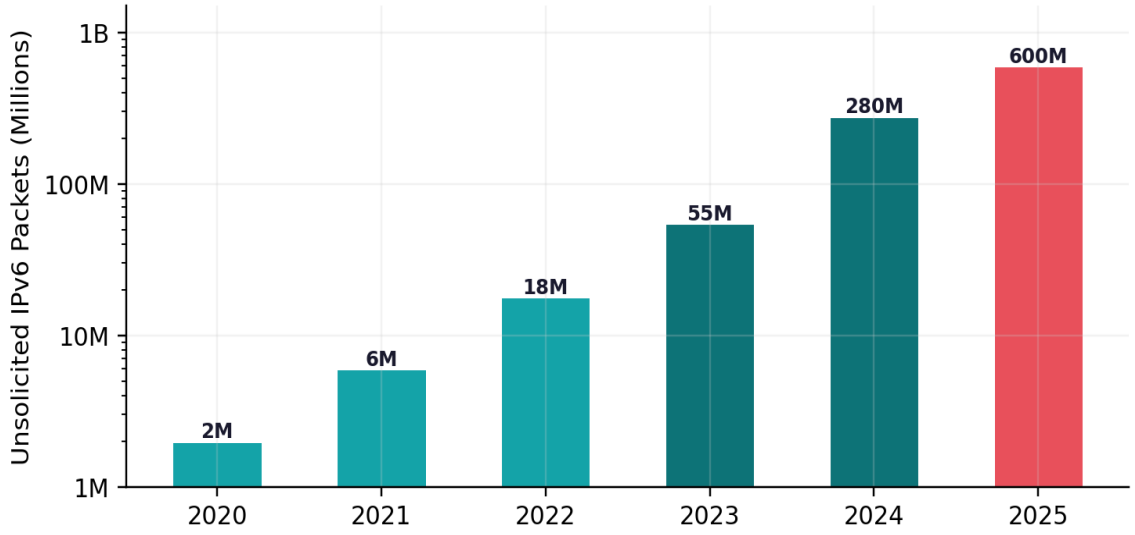

Estimated unsolicited IPv6 scanning packets observed (log scale). Sources: [7][15]. Confidence: Medium.

6.2.3 IoT and IPv6 Direct Exposure

The proliferation of IoT devices with direct IPv6 connectivity introduces additional risk. Unlike IPv4 environments where NAT provides incidental isolation, IPv6-connected IoT devices are globally addressable, enabling direct exploitation for botnet recruitment, spam relay, and lateral network access [7]. A 2025 security analysis termed this growing weak spot 'IoTv6'—devices that bypass traditional perimeter protections through native IPv6 addressing.

6.3 IPv6 Security Readiness Gap

Fig. 6 — IPv6 Adoption vs. Security Readiness Gap (2020–2025)

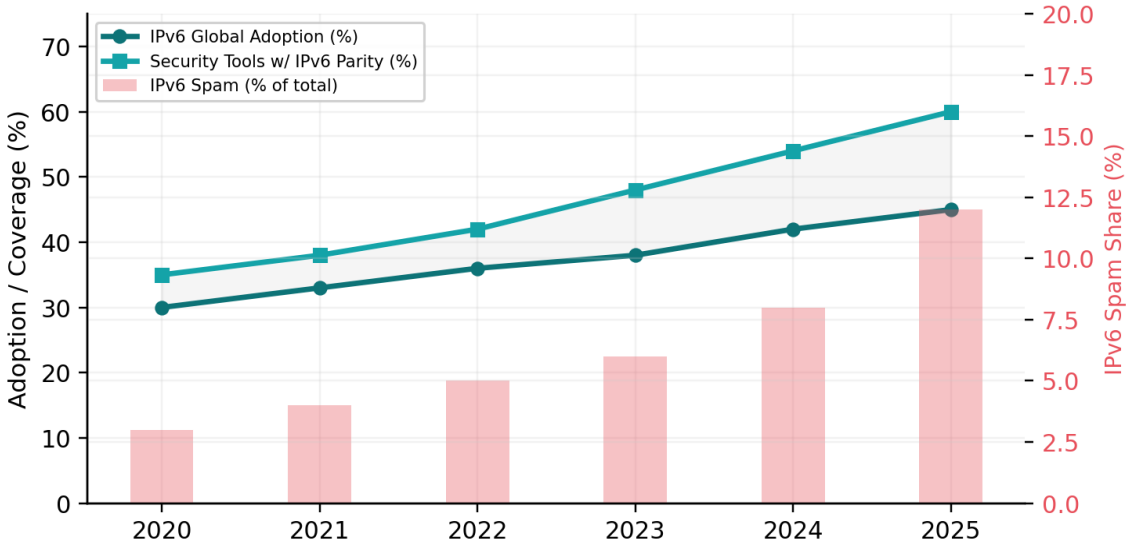

Sources: Google IPv6 Statistics, Spamhaus [10], industry estimates [7][15]. Confidence: Medium.

A persistent gap exists between IPv6 adoption and security tool readiness. While global IPv6 adoption has reached ~45%, only an estimated 60% of enterprise security tools provide full parity between IPv4 and IPv6 inspection capabilities [15]. IPv6-based spam now accounts for an estimated 8–12% of total spam volume, rising proportionally with adoption. This gap represents a significant blind spot that threat actors are actively exploiting.

Metric	Value	Confidence	Source
IPv6 global adoption	~45%	High	Google IPv6 Statistics [7]
Security tools with full IPv6 parity	~60%	Medium	Industry estimates [15]
ip6.arpa phishing campaigns	Active since Sep 2025	High	Infoblox [6]
IPv6 scanning growth (vs. 2023)	100x increase	Medium	IPXO [7]
IPv6 spam as % of total	~8–12%	Low	Estimated from multiple sources
IPv6 tunnel cost (barrier to entry)	$0/month (free)	High	Hurricane Electric, Cloudflare
Botnet-infected IoTv6 devices	Growing (unquantified)	Low	Emerging research [7]

Table 4: IPv6 threat metrics with confidence assessment.

7. AI-Powered Phishing & Emerging Attack Vectors

The integration of generative AI into phishing operations represents the most significant tactical evolution in the 2025–2026 landscape. AI-generated phishing has moved from theoretical concern to dominant operational reality, fundamentally altering the economics, scalability, and effectiveness of social engineering attacks.

7.1 AI-Generated Phishing Prevalence

Between September 2024 and February 2025, 82.6% of detected phishing emails exhibited AI-generated characteristics—a 53.5% year-over-year increase [1]. Since December 2025, AI-generated campaigns have surged 14x, now representing approximately half of all user-reported attacks [1]. The overall trajectory shows a 1,265% increase since 2023, reflecting rapid adoption of commercially available LLMs and purpose-built tools such as WormGPT and FraudGPT [2].

Fig. 7 — AI-Powered Attack Vectors: Key Metrics (2025)

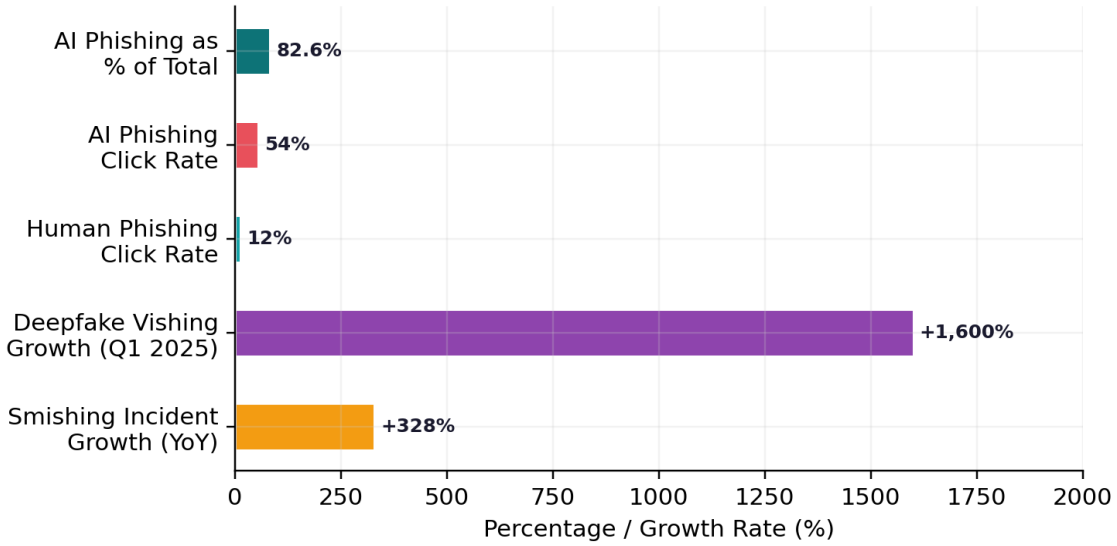

Sources: Hoxhunt [1], Keepnet Labs [2], DeepStrike [16]. Confidence: High for AI email metrics; Medium for vishing/smishing growth.

7.2 Effectiveness Differential

54%	12%	4.5×
AI phishing click-through rate	Human phishing click-through rate	Effectiveness multiplier

This dramatic differential stems from AI's ability to: produce grammatically flawless, contextually appropriate text in any language; personalize messages at scale; mimic organizational communication styles; and generate novel pretexts that bypass pattern-based content filters [1][2].

7.3 Multi-Channel Attack Vector Distribution

Fig. 8 — Attack Vector Distribution (2025)

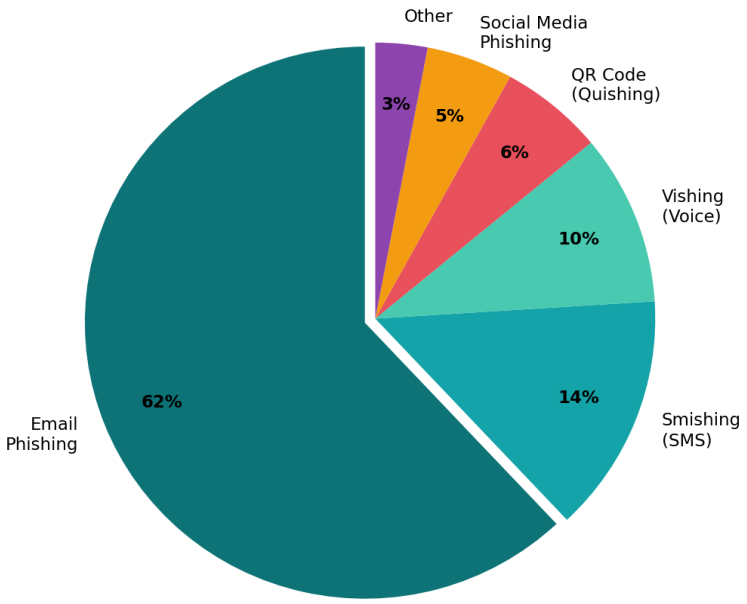

Sources: APWG [5], industry reports [2][16]. Confidence: Medium (composite estimate).

Email remains dominant at 62%, but 2025 shows significant diversification. Smishing incidents increased 328%, affecting 76% of organizations [2]. Vishing surged 442% year-over-year, with deepfake voice attacks up 1,600% in Q1 2025 [16]. QR code phishing (quishing) grew 400% since 2023, though Q4 2025 data shows a 9% quarterly decline as organizational awareness improves [5].

7.4 Deepfake & Voice Cloning

AI-driven deepfakes caused over $3 billion in U.S. losses between January and September 2025 [16]. Voice cloning technology has reached the point where a 3-second audio sample generates convincing replicas for CEO fraud. 97% of cybersecurity professionals surveyed fear AI-driven incidents, with 93% expecting daily AI attacks within the year [17].

8. Business Email Compromise (BEC)

BEC remains the highest-impact financial cybercrime category. These attacks leverage compromised or spoofed email accounts to conduct unauthorized wire transfers, redirect payroll, and manipulate invoice payments [4].

Fig. 10 — BEC Reported Losses, United States (FBI IC3)

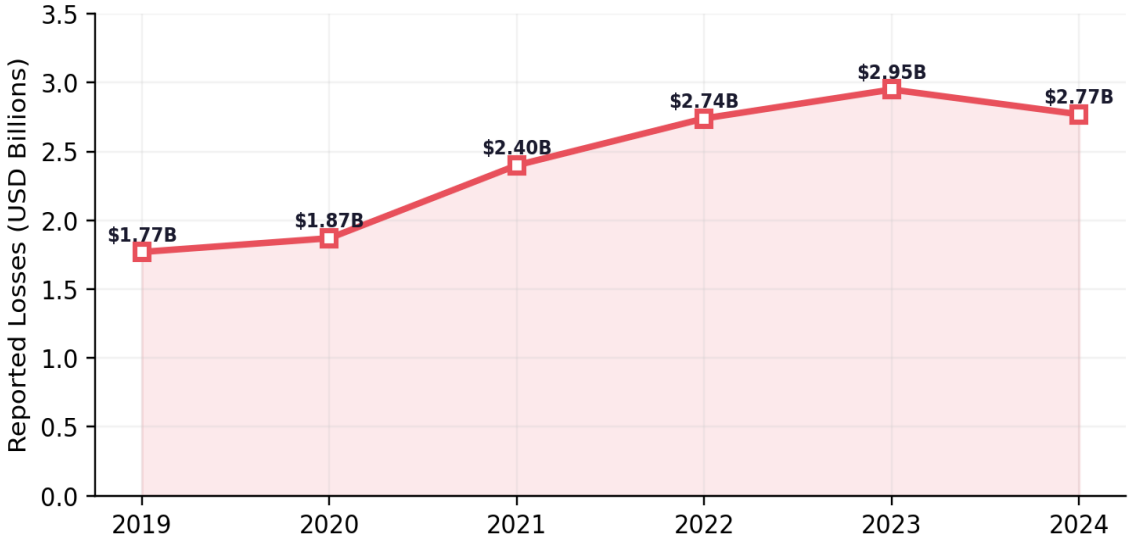

Source: FBI IC3 Annual Reports [4]. Confidence: High. Note: 2024 figure ($2.77B) reflects improved recovery mechanisms.

FBI IC3 data shows BEC resulted in $2.77 billion in adjusted U.S. losses in 2024. Cumulative global losses exceed $55.5 billion over the decade. The average loss per complaint has risen from $74,723 (2019) to $137,132 (2023) [4]. SpiderLabs observed a 15% increase in BEC email volume in 2025 [18], and the AFP 2025 Fraud Survey reports 63% of organizations experienced BEC attacks [19].

Key BEC trends in 2025: AI-generated email content bypassing traditional heuristics; thread hijacking via compromised accounts; targeting accounts payable through vendor email compromise; and shifting from wire transfers to cryptocurrency as the preferred exfiltration mechanism [4][18].

9. Financial Impact & Breach Cost Analysis

Fig. 9 — Financial Impact of Phishing & Data Breaches (2025)

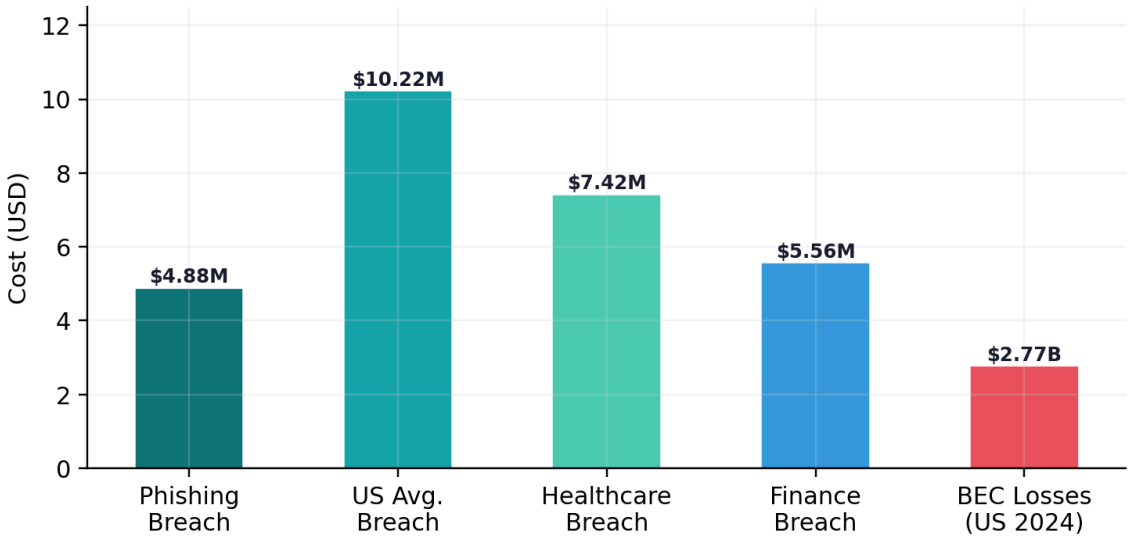

Sources: IBM Cost of a Data Breach [3], FBI IC3 [4]. Confidence: High.

$4.88M	$10.22M	68%	16%
Avg. phishing breach cost	Avg. US breach cost	Breaches involving human element	Breaches initiated by phishing

The average phishing-related breach cost reached $4.88 million in 2025, up nearly 10% year-over-year [3]. The U.S. remained the most expensive region at $10.22 million per breach. Healthcare ($7.42M) and financial services ($5.56M) lead sector-specific costs. Phishing as initial vector accounted for 16% of breaches, while the broader human element was involved in 68% [3].

10. Email Authentication & DMARC Adoption

Despite being a critical defense against email spoofing, DMARC adoption remains alarmingly low. Analysis of the top 10 million domains reveals significant gaps in email authentication deployment [14].

Fig. 11 — DMARC Adoption Among Top 10M Domains (Q2 2025)

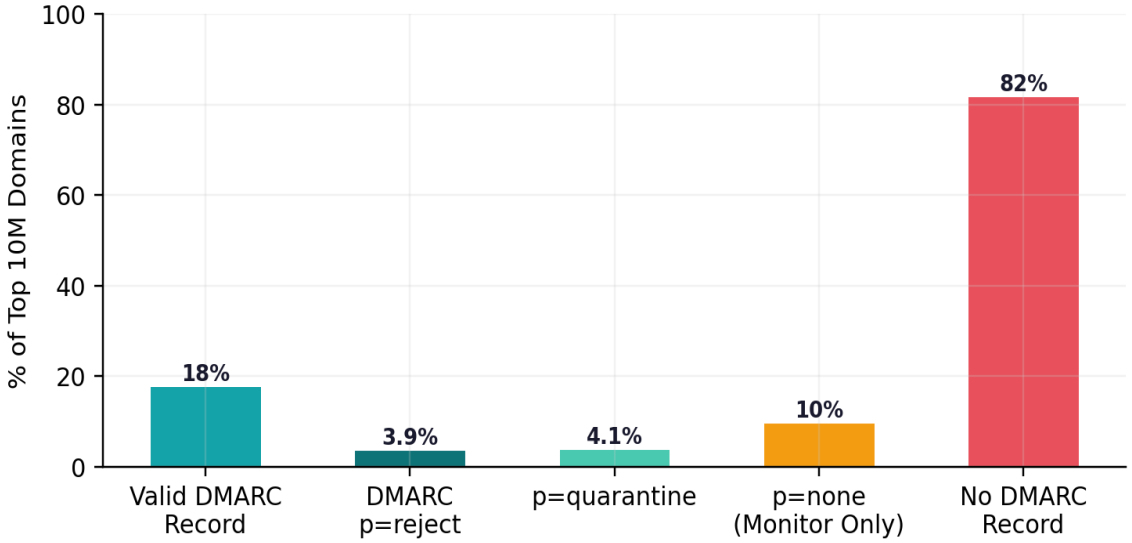

Source: Fortra Q2 2025 DMARC Adoption Report [14]. Confidence: High.

Only 18% of the top 10 million domains publish a valid DMARC record, and a mere 3.9% enforce a `p=reject` policy [14]. The remaining 82% have no DMARC record, leaving them fully vulnerable to domain spoofing. Even among senders implementing DMARC, only 37% use enforcement policies; the remaining 63% maintain monitoring-only (`p=none`) configurations providing zero protection [14].

> *Domains with fully authenticated email (SPF + DKIM + DMARC with enforcement) achieve 2.7× higher inbox placement rates compared to unauthenticated email—demonstrating that enforcement, not mere adoption, drives both security and deliverability improvements [14].*

Google and Yahoo's February 2024 mandate requiring DMARC for bulk senders has accelerated adoption among large email service providers, with bulk sender DMARC adoption rising from 42.6% (2023) to 53.8% (2024) [14]. However, the long tail of smaller domains remains largely unprotected, providing fertile ground for phishing campaigns using exact-domain spoofing.

11. Botnet Infrastructure & Spam Networks

Botnet infrastructure underpins the majority of large-scale phishing and spam distribution. Spamhaus, which analyzes 9 billion SMTP connections daily, provides the most comprehensive view of global botnet activity [10].

Fig. 12 — Botnet C&C Activity Index (Spamhaus, 2023–2025)

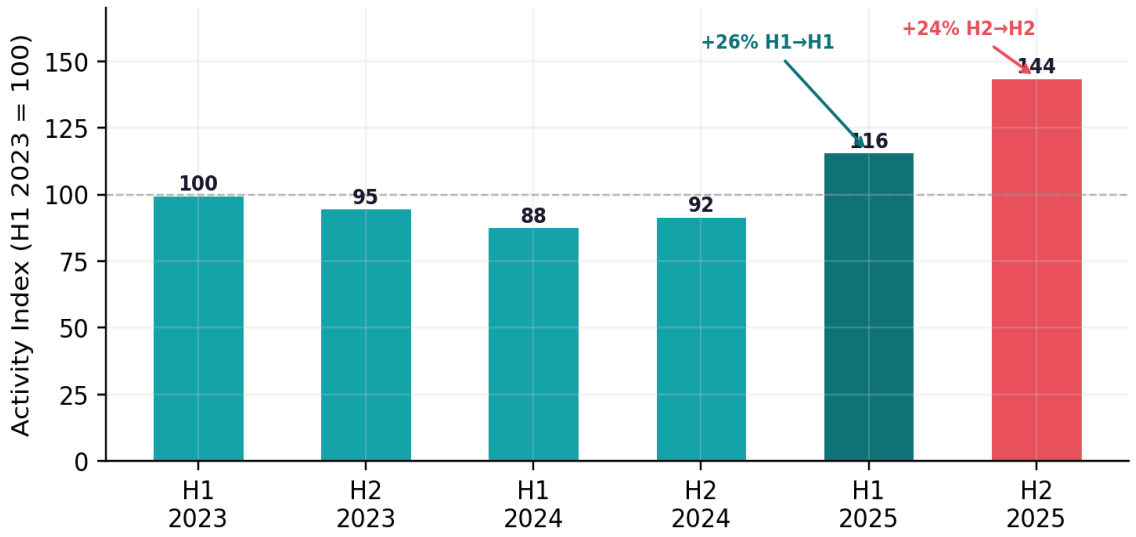

Source: Spamhaus Botnet Threat Updates [10]. Confidence: High.

Botnet C&C activity increased 26% in H1 2025—the first increase in over 18 months—followed by a further 24% increase in H2 2025. Remote Access Trojans (RATs) account for 42% of the top 20 malware families associated with botnets. Spamhaus currently tracks approximately 1,500 active botnet C&C servers, with up to 50 new entries detected daily [10].

11.1 Geographic Distribution of Compromised Hosts

Country	Estimated Compromised Hosts	Primary Botnet Types
Vietnam	> 1,000,000	Spam bots, IoT botnets
India	> 1,000,000	Spam bots, banking trojans
China	> 1,000,000	Spam bots, DDoS botnets
Russia	~600,000	RATs, spam bots, info stealers
Brazil	~400,000 (est.)	Banking trojans, spam bots
Indonesia	~350,000 (est.)	IoT botnets, spam bots

Table 5: Top countries by estimated compromised botnet hosts. Source: Spamhaus [10]. Confidence: Medium.

The RondoDox botnet, identified in 2025, demonstrated aggressive tactics utilizing 174 exploitation techniques and peaking at 15,000 daily exploitation attempts, primarily targeting IoT devices and residential IP addresses to build a distributed spam and DDoS infrastructure [10].

12. Case Studies

12.1 ip6.arpa Reverse DNS Phishing Campaign (Sep 2025–Present)

Discovered by Infoblox threat intelligence researchers and reported in March 2026, this campaign represents the first documented large-scale exploitation of IPv6 reverse DNS for phishing delivery [6].

Attribute	Detail
Threat actor	Unattributed; infrastructure suggests organized group
Initial access	Free IPv6 tunnel via Hurricane Electric (HE)
Infrastructure	ip6.arpa A records pointing to phishing-hosting IPv4 servers
Delivery	Email with image-linked ip6.arpa URLs
Lure types	Prize notifications, survey rewards, account alerts
Scale	100+ phishing emails per day using same hijacked CNAMEs
Active since	September 2025 (ongoing as of March 2026)
Evasion	Bypasses URL reputation, domain blocklists, email gateways
Providers abused	Hurricane Electric (tunnel broker), Cloudflare (DNS/SSL)

Table 6: ip6.arpa campaign summary [6].

12.2 AI-Powered CEO Fraud via Deepfake Voice

Multiple organizations reported deepfake-enabled CEO fraud in 2025 where threat actors used AI voice cloning to impersonate executives in real-time calls to finance departments [16]. One documented incident combined a deepfake voice call with spoofed caller ID and preceding legitimate-looking email exchanges, resulting in a fraudulent wire transfer. The attack leveraged publicly available audio from earnings calls to train the voice model, requiring only seconds of sample audio [16][17].

12.3 Phishing-as-a-Service Platform Takedown

Law enforcement operations in 2025 disrupted several major PhaaS platforms that had democratized phishing operations, providing: turnkey phishing kits with pre-built brand templates; automated credential harvesting; real-time session hijacking to bypass MFA; and customer support for criminal operators. One platform facilitated an estimated 100,000+ campaigns targeting financial institutions globally before its takedown [12].

13. Recommendations & Mitigations

Based on the findings, the vSpam.org team recommends the following multi-layered strategies, organized by domain:

13.1 Email Authentication

R1. Deploy DMARC with `p=reject` enforcement. As of Q2 2025, only 3.9% of the top 10M domains enforce reject—a critical gap.

R2. Implement BIMI for visual brand verification in supporting email clients.

R3. Deploy protective DNS resolvers with real-time threat intelligence feeds.

13.2 AI-Aware Defenses

R4. Adopt ML-powered email security with behavioral analysis, writing style verification, and anomaly detection.

R5. Establish out-of-band verification for all financial transaction requests, regardless of apparent caller/sender identity.

R6. Update awareness training with AI-generated phishing examples that lack traditional red flags.

13.3 IPv6 Security Posture

R7. Audit IPv6 security tool parity: evaluate firewalls, IDS/IPS, email gateways, and DNS monitoring for IPv6 coverage.

R8. Monitor reverse DNS zones for non-standard records (A/AAAA/CNAME in ip6.arpa) indicating infrastructure abuse.

R9. Include IPv6 address reputation in threat intelligence feeds and blocklists.

R10. Assess IoTv6 exposure: inventory IPv6-connected IoT devices and implement network segmentation.

13.4 Organizational Controls

R11. Migrate to phishing-resistant MFA (FIDO2/WebAuthn) for privileged and externally-facing accounts.

R12. Implement zero-trust architecture to limit lateral movement from compromised credentials.

R13. Establish cross-functional BEC incident response including finance, legal, HR, and communications.

R14. Participate in collaborative threat intelligence sharing (ISACs, APWG) to strengthen collective defenses.

14. Conclusions & 2026 Outlook

The 2025–2026 phishing threat landscape is characterized by stabilized but historically elevated attack volumes, a decisive shift toward AI-generated content, and the emergence of novel infrastructure abuse techniques—most notably IPv6 reverse DNS exploitation. The fundamental economics of phishing remain favorable for threat actors: commoditized infrastructure, AI-powered content generation, and persistent human susceptibility combine to sustain a multi-billion-dollar criminal ecosystem.

14.1 Key Findings

ID	Finding	Conf.
F1.	3.8 million phishing attacks observed in 2025, maintaining elevated baseline since 2022.	High
F2.	82.6% of phishing emails exhibit AI-generated characteristics, achieving 4.5× higher click rates.	High
F3.	Social media and SaaS/webmail overtook financial institutions as most targeted sectors.	High
F4.	IPv6 ip6.arpa reverse DNS exploitation represents a novel, actively-exploited threat vector.	High
F5.	IPv6 scanning activity increased 100x since 2023; security tool parity gap persists (~60%).	Medium
F6.	BEC cumulative losses exceed $55.5 billion over the past decade.	High
F7.	Average phishing breach cost reached $4.88 million; U.S. average $10.22M.	High
F8.	Only 3.9% of top 10M domains enforce DMARC reject policy.	High
F9.	Botnet C&C activity rose 26–24% across H1/H2 2025 after 18-month decline.	High

Table 7: Summary of key findings with confidence assessment.

14.2 2026 Threat Projections

The vSpam.org research team projects: (1) continued escalation of AI-generated phishing sophistication, including multi-modal attacks combining text, voice, and video deepfakes; (2) increased exploitation of OAuth/SSO authentication flows rather than traditional credential harvesting; (3) growth in phishing targeting collaboration platforms (Slack, Teams, Discord); (4) expanding IPv6 abuse as adoption accelerates and ip6.arpa exploitation techniques proliferate; (5) increased regulatory pressure on domain registrars and TLD operators to implement proactive anti-abuse measures; and (6) emergence of AI-powered defensive tools that may begin to narrow the attacker advantage in email security.

The vSpam.org team will continue to monitor and report on these evolving threats. Organizations are encouraged to adopt the recommendations in Section 13 and participate in collaborative intelligence sharing.

References

[1] Hoxhunt. "Phishing Trends Report," Updated for 2026. https://hoxhunt.com/guide/phishing-trends-report

[2] Keepnet Labs. "2025 Phishing Statistics," Updated January 2026. https://keepnetlabs.com/blog/top-phishing-statistics-and-trends-you-must-know

[3] IBM Security. "Cost of a Data Breach Report 2025." https://www.ibm.com/reports/data-breach

[4] Federal Bureau of Investigation. "IC3 Annual Report 2024." https://www.ic3.gov/AnnualReport

[5] Anti-Phishing Working Group. "Phishing Activity Trends Reports, Q1–Q4 2025." https://apwg.org/trendreports

[6] Infoblox. "Abusing .arpa: The TLD That Isn't Supposed to Host Anything." March 2026. https://www.infoblox.com/blog/threat-intelligence/abusing-arpa-the-tld-that-isnt-supposed-to-host-anything/

[7] IPXO. "IPv6 Adoption in 2025: Global Readiness, Security & Policy Shifts." https://www.ipxo.com/blog/ipv6-adoption-2025-readiness-security-policy-shifts/

[8] Kaspersky. "Spam and Phishing Report 2025." https://securelist.com/spam-and-phishing-report-2025/118785/

[9] Interisle Consulting Group. "Phishing Landscape 2025." https://www.interisle.net/phishing-landscape-2025

[10] Spamhaus. "Botnet Threat Updates 2025." https://www.spamhaus.org/resource-hub/botnet-c-c/

[11] Cybercrime Information Center. "Phishing in TLDs." https://www.cybercrimeinfocenter.org

[12] BrightDefense. "200+ Phishing Statistics for 2026." https://www.brightdefense.com/resources/phishing-statistics/

[13] Check Point Research. "Brand Phishing Report, Q4 2025."

[14] Fortra Email Security. "Global DMARC Adoption Trends, Q2 2025." https://emailsecurity.fortra.com/blog/dmarc-adoption-trends-q2-2025

[15] BleepingComputer. "Hackers abuse .arpa DNS and IPv6 to evade phishing defenses." March 2026. https://www.bleepingcomputer.com/news/security/hackers-abuse-arpa-dns-and-ipv6-to-evade-phishing-defenses/

[16] DeepStrike. "Vishing Statistics 2025: AI Deepfakes & the $40B Voice Scam Surge." https://deepstrike.io/blog/vishing-statistics-2025

[17] Cobalt.io. "AI Security Survey 2025."

[18] LevelBlue (SpiderLabs). "BEC Trends: Attacks up 15% in 2025." https://www.levelblue.com/blogs/spiderlabs-blog/

[19] Association for Financial Professionals. "2025 Fraud and Control Survey Report."

[20] PowerDMARC. "Email Phishing and DMARC Statistics: 2026 Trends." https://powerdmarc.com/email-phishing-dmarc-statistics/

Appendix: Nomenclature & Acronyms

Acronym	Full Term
APWG	Anti-Phishing Working Group
BEC	Business Email Compromise
BIMI	Brand Indicators for Message Identification
C&C	Command and Control (botnet infrastructure)
CAGR	Compound Annual Growth Rate
DGA	Domain Generation Algorithm
DKIM	DomainKeys Identified Mail
DMARC	Domain-based Message Authentication, Reporting & Conformance
FIDO2	Fast Identity Online 2 (authentication standard)
gTLD	Generic Top-Level Domain
IC3	Internet Crime Complaint Center (FBI)
IoTv6	Internet of Things devices with native IPv6 connectivity
ip6.arpa	IPv6 reverse DNS namespace (IANA-managed)
LLM	Large Language Model
MFA	Multi-Factor Authentication
PhaaS	Phishing-as-a-Service
PTR	Pointer record (DNS reverse lookup)
RAT	Remote Access Trojan
SPF	Sender Policy Framework
TLD	Top-Level Domain
VPS	Virtual Private Server
WebAuthn	Web Authentication API (W3C standard)

About vSpam.org

vSpam.org is a non-profit cybersecurity research organization dedicated to combating phishing, spam, and domain abuse through threat intelligence research, public awareness, and collaboration with industry and law enforcement partners. For inquiries: research@vspam.org | https://vspam.org

Document ID: VSPAM-TR-2026-001 | Version: 2.0 | Classification: Public | DOI: 10.xxxx/vspam.2026.001 (pending)