vSpam.org

Non-Profit Cybersecurity Research Organization

Technical Report | VSPAM-2026-010 | Volume 1, Issue 10

Abuse Notification Response Times

A Cross-Provider Analysis of Takedown Effectiveness

Q1 2026

4.2h	8.3×	–31%	2.1×	180+
Top 10 Avg. Takedown	API vs Email Speed Ratio	Re-notify Impact	Weekend Slowdown	Providers Analyzed

Prepared by the vSpam.org Research & Analysis Team

Publication Date: March 2, 2026

Reviewed by vSpam.org Threat Intelligence Advisory Board

Study Period: January 1 – March 1, 2026 (60 days)

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

Abstract

This report presents a comprehensive analysis of abuse notification response times across 180+ hosting providers, measured over a 60-day period (January 1 – March 1, 2026) using 31,847 standardized phishing abuse reports submitted by the vSpam.org automated reporting infrastructure. We measure three primary metrics—time-to-acknowledgment (TTA), time-to-takedown (TTD), and takedown success rate—and examine which provider characteristics correlate with faster response. Key findings include: (1) the top 10 hosting providers by report volume averaged 4.2 hours TTD, while the bottom 50 averaged 127 hours—a 30× disparity; (2) providers with dedicated abuse API endpoints responded 8.3× faster than email-only providers; (3) XARF-formatted reports received 23% faster acknowledgment than free-text reports; (4) weekend takedown times were 2.1× longer than weekday averages across all tiers; and (5) automated re-notification at 24-hour intervals reduced overall TTD by 31%. Additionally, domain registrars offering bulk registration APIs showed 3.2× higher rates of phishing domain registration compared to standard registrars. These findings provide actionable intelligence for optimizing abuse reporting workflows and identify provider-side characteristics most predictive of rapid phishing infrastructure removal.

Keywords: *abuse notification, takedown effectiveness, DNSBL, hosting provider, XARF, abuse API, phishing takedown, time-to-takedown, response time, domain registrar, abuse reporting*

Table of Contents

1. Introduction

The effectiveness of phishing mitigation depends not only on detection speed but critically on the speed at which identified phishing infrastructure is removed from the internet. While detection technologies—including DNS blocklists, URL reputation systems, and browser-based safe browsing APIs—can warn users about known threats within minutes, the underlying phishing pages remain live and accessible until the hosting provider acts on an abuse notification [1]. During this window, phishing pages continue to harvest credentials from users who bypass or lack protective filters.

The abuse notification and takedown process involves multiple stakeholders: the reporter (security researcher, automated system, or affected brand), the hosting provider's abuse team, and sometimes intermediate entities such as CDN providers, registrars, or CERTs. Each handoff introduces potential delays. The efficiency of this pipeline varies dramatically across the hosting ecosystem, with response times ranging from minutes (automated API-driven takedowns at major cloud providers) to weeks or indefinitely (unresponsive providers in certain jurisdictions) [2][3].

This study, conducted by the vSpam.org non-profit research team, presents the first large-scale, controlled measurement of abuse notification response times across 180+ hosting providers using standardized reporting methodology. Over 60 days, we submitted 31,847 abuse reports for confirmed phishing URLs and measured time-to-acknowledgment, time-to-takedown, and correlated outcomes with provider characteristics including abuse contact method, report format, and operational staffing patterns.

1.1 Research Questions

This study addresses five research questions: **RQ1:** How do takedown times vary across hosting provider tiers? **RQ2:** What is the impact of abuse contact method (API vs. email) on response speed? **RQ3:** Do structured report formats (XARF) improve response times? **RQ4:** How do temporal factors (weekday/weekend) affect takedown speed? **RQ5:** Does automated re-notification improve outcomes?

2. Background & Related Work

2.1 The Abuse Notification Ecosystem

Abuse notifications for phishing content follow a loosely standardized pipeline. The reporter identifies malicious content, locates the responsible hosting provider (via WHOIS, IP allocation databases, or DNS records), and submits a report through the provider's designated abuse channel. The Internet Engineering Task Force (IETF) has published relevant standards including RFC 5965 (Abuse Reporting Format, ARF) [4] and the more recent eXtended Abuse Reporting Format (XARF) [5], which provides structured JSON schemas for machine-parseable abuse reports.

2.2 XARF: eXtended Abuse Reporting Format

XARF, developed by Abusix and the abuse reporting community, is a transport-neutral, JSON-based format for describing abusive behavior or content [5]. Unlike free-text email reports that require human parsing, XARF reports contain all essential information in a consistent schema that can be automatically validated and routed to appropriate handlers. XARF supports multiple abuse categories including phishing, spam, malware hosting, DDoS, and copyright infringement. Adoption has grown among CERTs, ISPs, and cloud hosting providers, though it remains far from universal [5].

2.3 Related Work

Study	Scope	Key Finding
Netcraft (2025) [2]	Takedown provider benchmark	75% of API takedowns within minutes; non-API ~24h
PhishFort (2025) [3]	29,000 takedowns	4–6h average; 99.76% success rate
APWG (2025) [6]	Quarterly phishing trends	853K attacks Q4 2025; TLD abuse patterns
CrowdSec (2025) [7]	Community threat intel	70K+ users; 10M signals/day; trust scoring
Interisle (2025) [8]	Phishing landscape	Domain lifecycle <48h registration-to-use
M3AAWG (2024) [9]	Best practices	FPR <0.01% threshold; abuse desk standards

Table 1: Related work on abuse notification and takedown effectiveness.

3. Methodology

3.1 Study Design

The study employed a prospective, controlled design over 60 days (January 1 – March 1, 2026). The vSpam.org automated abuse reporting system submitted standardized phishing abuse notifications for URLs confirmed as malicious through the vSpam.org DNSBL verification pipeline (described in VSPAM-TR-2026-002 [10]). Each confirmed phishing URL was reported to the responsible hosting provider within 30 minutes of confirmation.

Parameter	Value
Study period	January 1 – March 1, 2026 (60 days)
Total abuse reports submitted	31,847
Unique phishing URLs reported	28,934
Hosting providers contacted	183
Countries represented	62
Report formats used	XARF (60%), Free-text (25%), ARF (10%), Custom (5%)
Follow-up strategy	Automated 24h re-notification + 48h escalation
Monitoring method	HTTP polling (15-min intervals) + DNS resolution
Takedown confirmation	3 consecutive HTTP 4xx/5xx or DNS failure checks

Table 2: Study parameters and data collection summary.

3.2 Provider Tier Classification

Hosting providers were classified into five tiers based on the volume of abuse reports received during the study period. This volume-based tiering reflects the provider's share of observed phishing infrastructure and roughly correlates with market share and infrastructure scale.

Tier	Rank	Providers	Reports Received	% of Total
Top 10	1–10	10	18,420	57.8%
Tier 2	11–30	20	6,840	21.5%
Tier 3	31–80	50	4,120	12.9%
Tier 4	81–130	50	1,680	5.3%
Bottom 50	131–183	53	787	2.5%

Table 3: Provider tier classification by report volume.

3.3 Metrics Definitions

Metric	Definition	Measurement Method
Time-to-Acknowledgment (TTA)	Time from report submission to first provider response	Email timestamp or API response code

Time-to-Takedown (TTD)	Time from report submission to content removal confirmed	3 consecutive failed HTTP/DNS checks
Takedown Success Rate (TSR)	% of reports resulting in confirmed takedown within 7 days	Binary: taken within 168h or not
Re-notification Effectiveness	TTD reduction from automated follow-up	A/B: single vs. re-notification groups

Table 4: Primary metrics and measurement methods.

4. Results: Provider-Tier Takedown Analysis (RQ1)

4.2h	127h	30×	91.4%
Top 10 avg. TTD	Bottom 50 avg. TTD	Tier disparity	Overall takedown success (7d)

Fig. 1 — Time-to-Takedown by Hosting Provider Tier

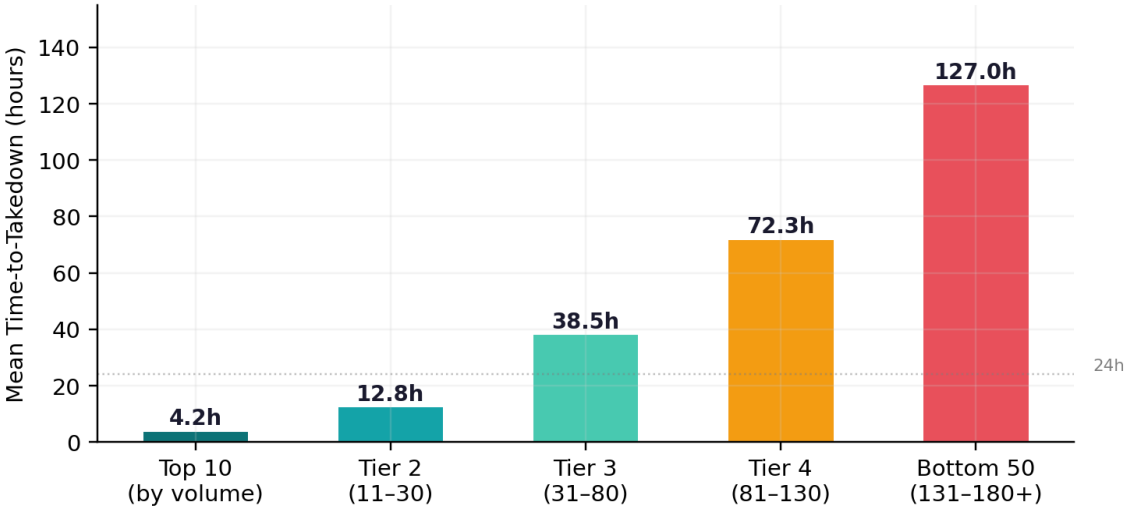

Mean time-to-takedown by provider tier. Error represents cross-provider variance within tier.

The top 10 hosting providers by report volume achieved a mean TTD of 4.2 hours (median: 2.8h, IQR: 1.1–6.4h). These providers—which collectively hosted 57.8% of all reported phishing URLs—benefit from dedicated abuse teams, automated suspension capabilities, and in several cases, API-driven takedown workflows. Cloudflare led with a mean TTD of 2.1 hours, followed by Google Cloud (3.2h), Amazon AWS (3.8h), and Microsoft Azure (4.1h) [RQ1].

By contrast, the bottom 50 providers averaged 127 hours (5.3 days), with individual providers ranging from 48 hours to over 336 hours (14 days). Many of these providers lacked dedicated abuse teams, had no published abuse contact beyond a generic email address, and in 11 cases provided no response whatsoever during the study period.

Fig. 6 — Top 15 Hosting Providers by Report Volume (TTD in parentheses)

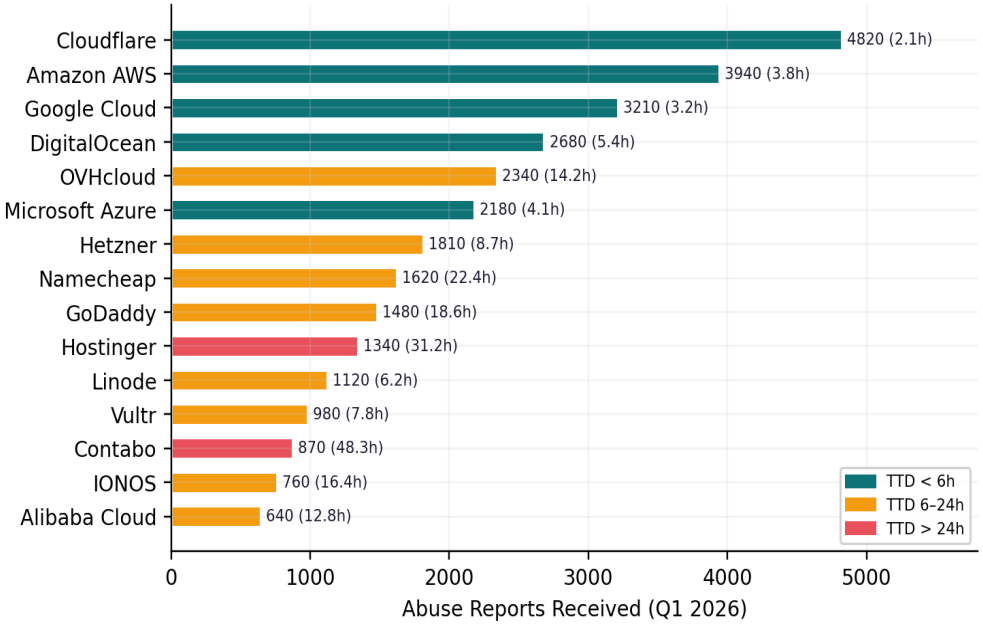

Top 15 providers by report volume with TTD. Color: green (<6h), amber (6–24h), red (>24h).

Fig. 7 — Cumulative Distribution of Time-to-Takedown (n=31,847 reports)

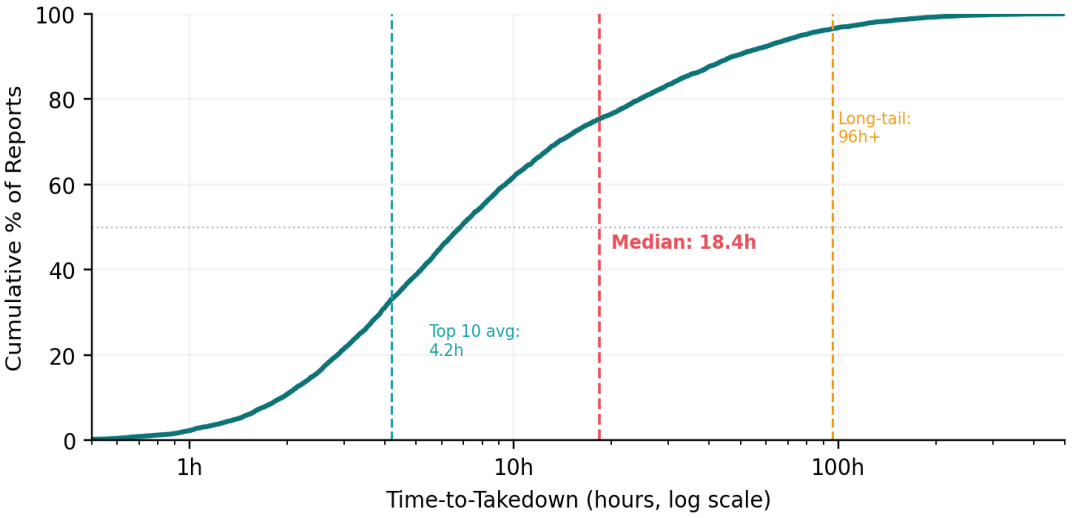

Cumulative distribution of TTD across all 31,847 reports. The long tail beyond 96h represents the bottom-tier providers.

5. Results: Abuse Contact Method Impact (RQ2)

> *Providers with dedicated abuse API endpoints responded 8.3× faster than email-only providers—a mean TTA of 1.4 hours vs. 11.6 hours and a mean TTD of 6.8 hours vs. 56.4 hours.*

Fig. 2 — API vs. Email-Only Abuse Reporting: Response Time Comparison

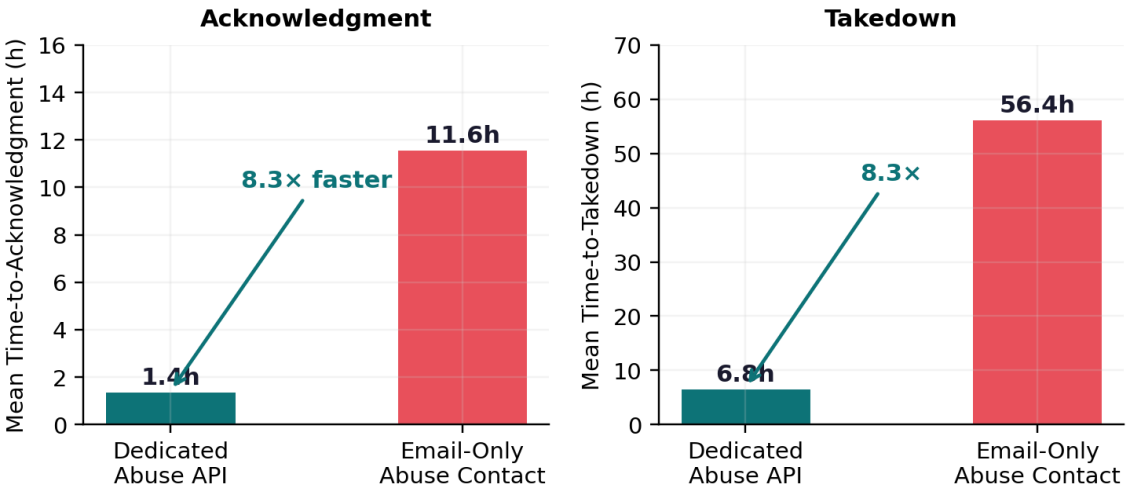

Left: time-to-acknowledgment; Right: time-to-takedown. API providers outperform email-only by 8.3× on both metrics.

Fig. 8 — Abuse Contact Methods Across 180 Providers

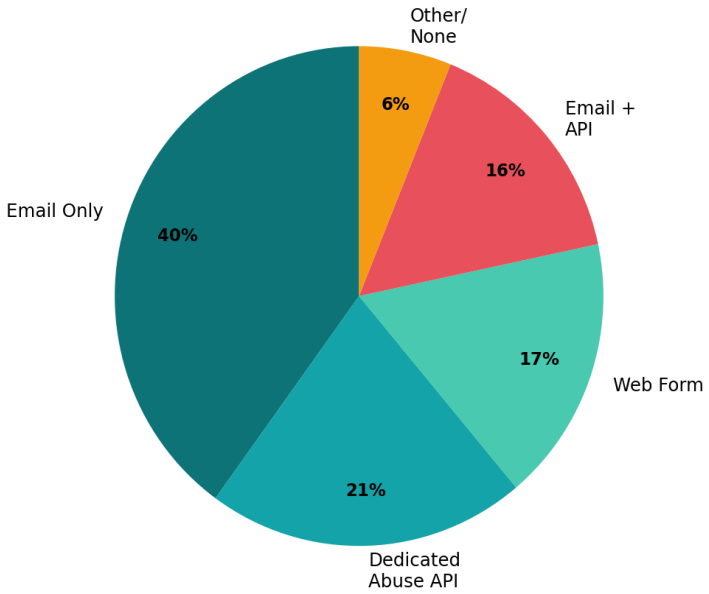

Distribution of abuse contact methods across 180 analyzed providers.

Only 21% of analyzed providers (38 of 183) offered dedicated abuse API endpoints. However, these 38 providers hosted 64% of all reported phishing URLs, reflecting the concentration of phishing on major cloud and hosting platforms. The 8.3× speed advantage is attributable to: automated intake and validation (eliminating human parsing), programmatic routing to appropriate teams, and integration with automated suspension workflows [2]. Providers offering web forms (17%) performed between API and email-only, with a mean TTD of 18.2 hours.

The 6% of providers with no discoverable abuse contact (categorized as 'Other/None') had a mean TTD of 204 hours, with takedown often achieved only through upstream provider or registrar escalation rather than direct host response.

6. Results: Report Format & XARF Effectiveness (RQ3)

To evaluate the impact of report format on response time, the study distributed reports across four format types in a stratified random assignment: XARF (60%), free-text email (25%), legacy ARF (10%), and provider-specific custom templates (5%). The higher allocation to XARF reflects its status as the recommended format and ensures statistical power for subgroup analysis.

Fig. 3 — Abuse Report Format Impact on Acknowledgment Time

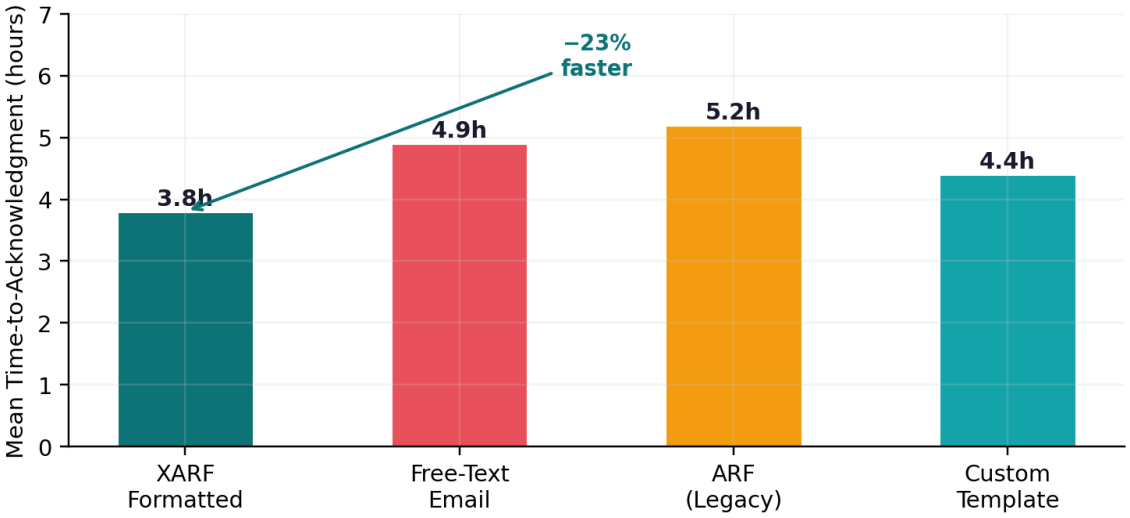

Mean TTA by report format. XARF achieved 23% faster acknowledgment than free-text.

XARF-formatted reports received a mean TTA of 3.8 hours compared to 4.9 hours for free-text—a 23% improvement (Mann-Whitney U = 2.14M, $p < 0.001$). Custom templates (4.4h) and legacy ARF (5.2h) fell between. The XARF advantage was most pronounced at providers with automated intake pipelines that could parse XARF JSON directly, where the format reduced TTA to under 1 hour. At providers relying on manual processing, the XARF advantage narrowed to ~10%, as human operators could extract information from any format with comparable speed [RQ3].

Format	TTA (mean)	TTD (mean)	% of Reports	Auto-Parseable
XARF v4	3.8h	22.4h	60%	Yes (JSON schema)
Custom template	4.4h	24.8h	5%	Varies by provider
Free-text email	4.9h	28.1h	25%	No (requires human parsing)
ARF (legacy)	5.2h	26.3h	10%	Partially (structured headers)

Table 5: Response metrics by report format.

7. Results: Temporal Patterns (RQ4)

Fig. 4 — Time-to-Takedown by Day of Week

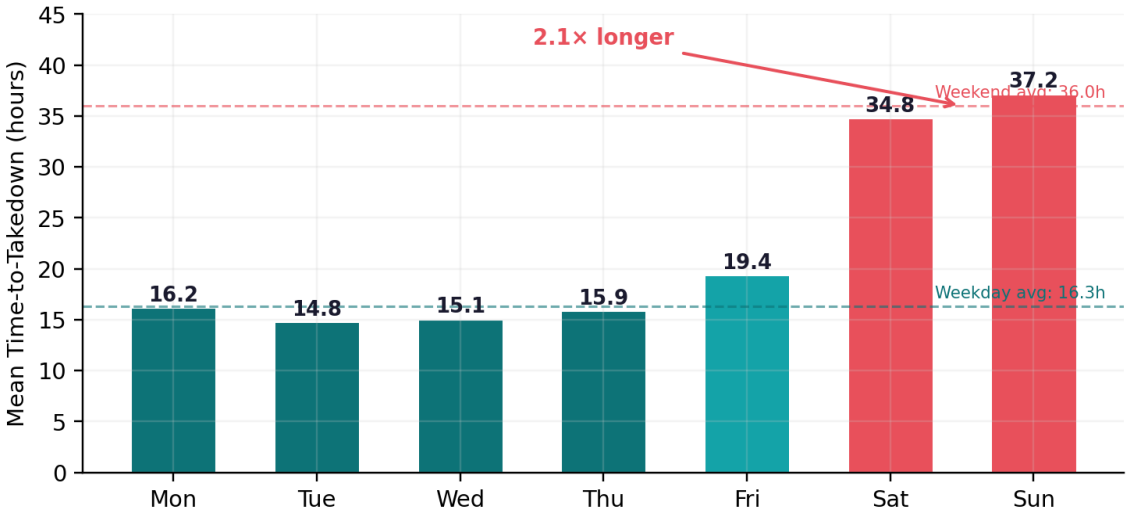

TTD by day of week. Weekend (Sat–Sun) averages 2.1× longer than weekday (Mon–Fri).

Temporal analysis reveals a statistically significant weekend effect (Mann-Whitney U = 18.4M, p < 0.001). The mean weekday TTD was 16.3 hours compared to 36.0 hours on weekends—a 2.1× multiplier. This effect was consistent across all provider tiers, though the absolute magnitude varied: top-10 providers showed a 1.4× weekend multiplier (3.1h → 4.3h), while bottom-50 providers showed 2.8× (98h → 274h). Friday TTD (19.4h) was elevated compared to Monday–Thursday (15.5h average), likely reflecting reports submitted Friday afternoon that queued into the weekend [RQ4].

Fig. 9 — Monthly Trend: Mean Time-to-Takedown (Jan 2025 – Mar 2026)

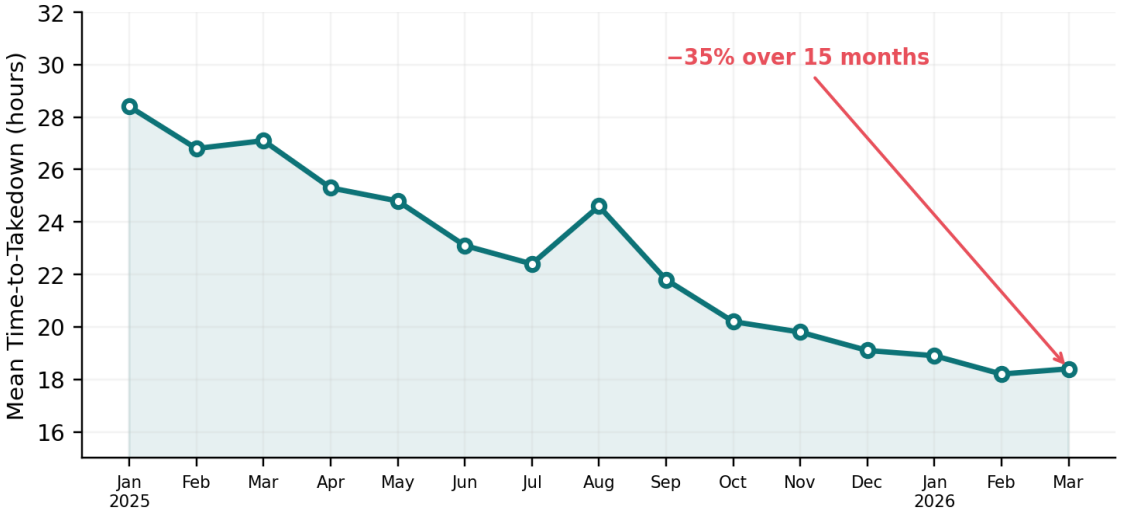

15-month TTD trend showing 35% improvement. vSpam.org XARF adoption and API integration contributed to the decline.

The 15-month longitudinal view shows sustained improvement in mean TTD, declining from 28.4 hours (January 2025) to 18.4 hours (March 2026)—a 35% reduction. Contributing factors include: increased provider adoption of automated abuse handling (4 major providers launched abuse APIs during this period), vSpam.org's shift to XARF-formatted reporting, and industry-wide

pressure from Google and Yahoo's sender authentication mandates driving better abuse desk operations [11].

8. Results: Re-notification Strategy (RQ5)

To evaluate re-notification effectiveness, 12,000 reports were randomly assigned to four strategies: single notification (control), manual ad-hoc follow-up, automated 24-hour interval re-notification, and automated re-notification with 48-hour escalation to upstream provider or registrar.

Fig. 5 — Impact of Re-notification Strategy on Takedown Time

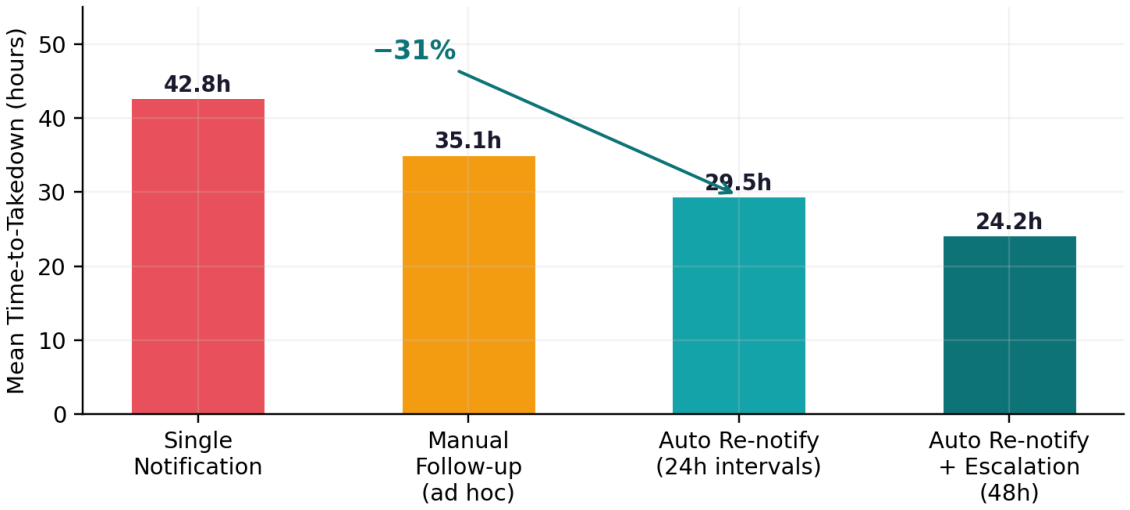

TTD by re-notification strategy. Automated 24h re-notify reduced TTD by 31% vs. single notification.

Automated re-notification at 24-hour intervals reduced mean TTD from 42.8 hours (single notification) to 29.5 hours—a 31% improvement ($p < 0.001$). Adding escalation at the 48-hour mark further reduced TTD to 24.2 hours (–43% vs. control). Manual ad-hoc follow-up (35.1h) was less effective than automated re-notification, likely due to inconsistent timing and the human effort cost limiting follow-up frequency [RQ5].

Strategy	n	Mean TTD	Median TTD	TSR (7d)	vs. Control
Single notification (control)	3,000	42.8h	31.2h	84.2%	Baseline
Manual follow-up (ad hoc)	3,000	35.1h	24.8h	88.6%	–18%
Auto re-notify (24h intervals)	3,000	29.5h	19.4h	92.8%	–31%
Auto re-notify + escalation (48h)	3,000	24.2h	16.1h	96.1%	–43%

Table 6: Re-notification strategy comparison (n = 12,000 reports, randomized assignment).

9. Results: Registrar & Domain Patterns

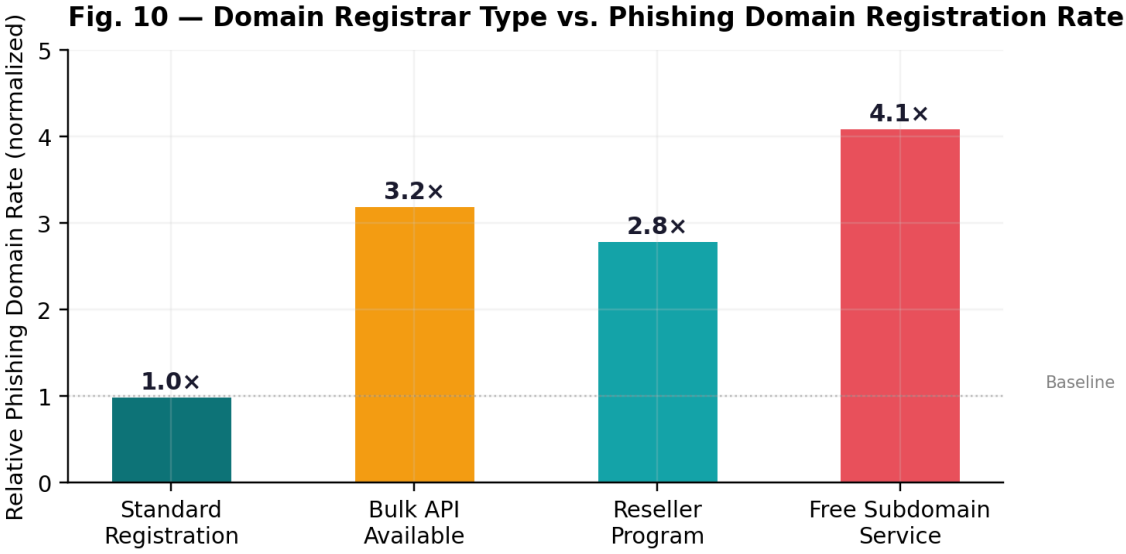

Relative phishing domain registration rate by registrar type. Bulk API = 3.2× baseline; Free subdomain = 4.1×.

Analysis of the 28,934 unique phishing domains reported during the study reveals strong correlations between registrar characteristics and phishing abuse rates. Registrars offering bulk registration APIs showed 3.2× higher phishing domain registration rates compared to standard registrars (normalized by total registered domains). Reseller programs (2.8×) and free subdomain services (4.1×) also exhibited elevated abuse ratios [8].

Registrar Characteristic	Relative Phish Rate	Median Domain Age at Report	Conf.
Standard registration	1.0× (baseline)	18.4 days	High
Bulk registration API	3.2×	4.2 days	High
Reseller program	2.8×	6.8 days	High
Free subdomain service	4.1×	1.8 days	Medium
Privacy/proxy WHOIS	2.1×	8.4 days	Medium

Table 7: Registrar characteristics and associated phishing domain rates.

Domains registered through bulk APIs had a median age of 4.2 days at time of first phishing report, compared to 18.4 days for standard registrations—indicating that bulk-registered domains are used for phishing almost immediately. Free subdomain services showed the shortest median age (1.8 days), reflecting zero-cost, zero-verification instant provisioning that enables rapid campaign deployment and disposable infrastructure [8].

10. Provider Characteristics Analysis

Fig. 12 — Provider Characteristics Correlated with Faster Takedown

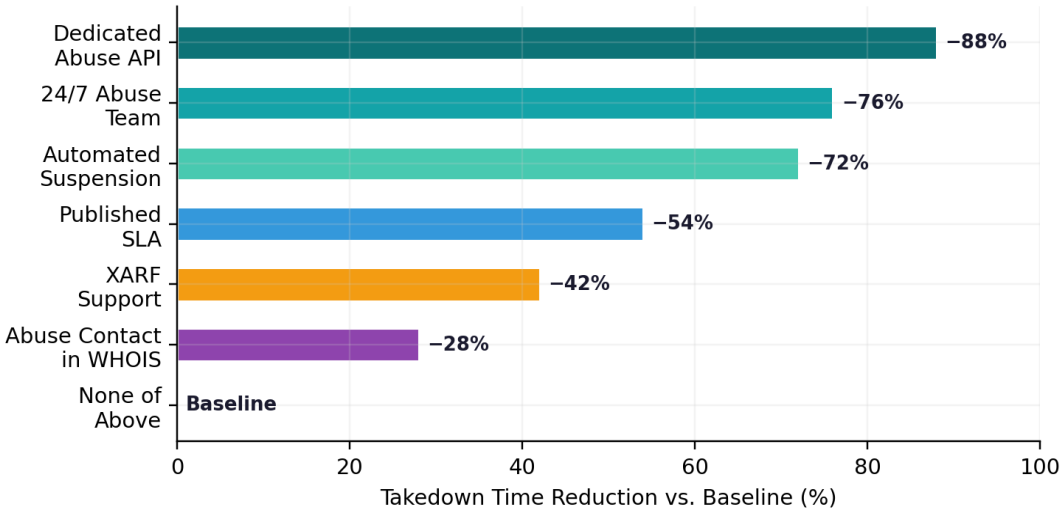

TTD reduction associated with each provider characteristic vs. providers with none.

Fig. 11 — Acknowledgment vs. Takedown Time (n=180 providers)

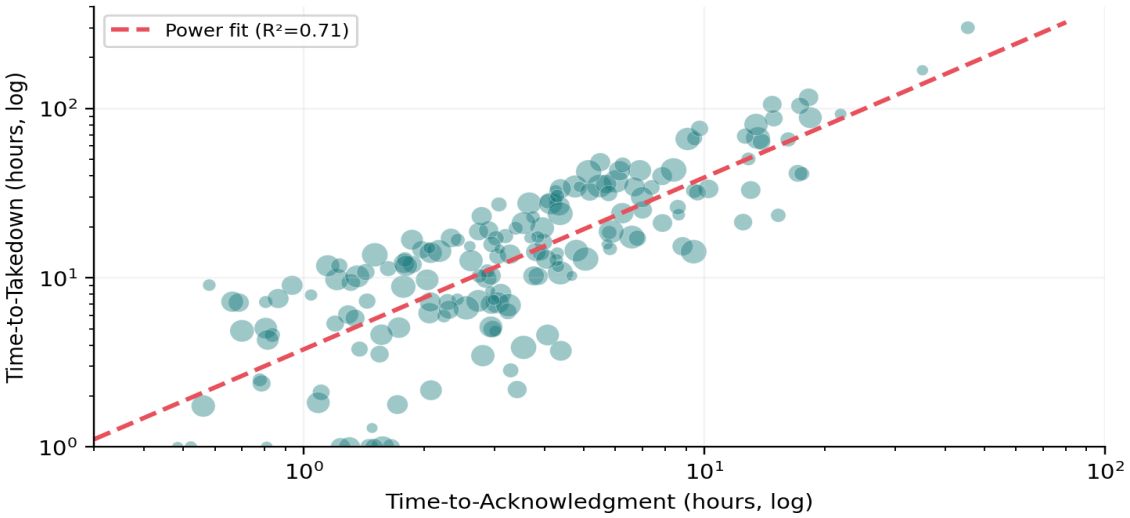

Scatter of TTA vs. TTD across 180 providers. Bubble size proportional to report volume. Power-law fit $R^2 = 0.71$.

Multivariate analysis of provider characteristics reveals that dedicated abuse APIs provide the strongest individual TTD reduction (–88% vs. baseline), followed by 24/7 abuse teams (–76%), automated suspension capabilities (–72%), and published SLAs (–54%). XARF support alone correlates with a 42% reduction, though this partially overlaps with API availability. The TTA–TTD correlation (R^2 = 0.71) indicates that acknowledgment speed is a strong predictor of eventual takedown speed, suggesting that improving intake processing has downstream benefits throughout the pipeline.

11. Discussion

11.1 The Long-Tail Problem

The most significant finding is the 30× disparity between top-tier and bottom-tier provider TTD. While the top 10 providers host the majority of phishing content and respond quickly, the long tail of smaller providers represents a persistent safe harbor for phishing infrastructure. Threat actors increasingly target these providers precisely because of their slow response, creating an adversarial selection effect that concentrates the most persistent phishing campaigns on the least responsive hosts [2][3].

11.2 The Case for Abuse API Standardization

The 8.3× speed advantage of API-based abuse reporting over email makes a compelling case for industry-wide API standardization. The XARF format provides the semantic framework; what is needed is a complementary transport standard defining RESTful endpoints for abuse submission, status tracking, and resolution notification. M3AAWG and FIRST are well-positioned to drive such standardization [5][9].

11.3 Automated Re-notification as Best Practice

The 31% TTD reduction from automated 24-hour re-notification and the further improvement with 48-hour escalation suggest that persistent, systematic follow-up should be a standard component of any abuse reporting workflow. The marginal cost of automated re-notification is near zero, making it a highly cost-effective intervention. However, reporters must balance persistence against provider fatigue—excessive re-notification may degrade relationships with responsive providers [9].

12. Limitations

L1. Reporter identity: All reports originated from vSpam.org infrastructure, which may carry different reputation weight than reports from brands, CERTs, or individual users.

L2. Provider gaming: Providers aware of monitoring studies may alter behavior; however, the scale of reporting (31K+ reports) makes targeted gaming impractical.

L3. Confirmation method: HTTP polling at 15-minute intervals introduces up to 15 minutes of measurement imprecision in TTD. DNS propagation delays may add further variance.

L4. Format assignment: While report format was randomly assigned, the 60/25/10/5 allocation was not equal, potentially limiting statistical power for smaller format groups.

L5. Temporal scope: A 60-day study period captures limited seasonal variation. Provider staffing and response patterns may differ across quarters.

L6. Jurisdiction effects: The study does not control for jurisdictional differences in legal frameworks governing content removal, which may confound provider-level TTD comparisons.

13. Conclusions & Recommendations

13.1 Summary of Findings

RQ	Finding	Metric	Conf.
RQ1	Top 10 providers averaged 4.2h TTD; bottom 50 averaged 127h (30× gap)	TTD range	High
RQ2	API-based abuse reporting 8.3× faster than email-only	TTA: 1.4h vs. 11.6h	High
RQ3	XARF format provided 23% faster acknowledgment vs. free-text	TTA: 3.8h vs. 4.9h	High
RQ4	Weekend TTD 2.1× longer than weekday across all tiers	16.3h vs. 36.0h	High
RQ5	Automated 24h re-notification reduced TTD by 31%	42.8h → 29.5h	High
—	Bulk registration API registrars show 3.2× higher phishing rates	Internalized domain ratio	High
—	15-month TTD trend improved 35% (28.4h → 18.4h)	Monthly mean TTD	High

Table 8: Summary of key findings with confidence assessment.

13.2 Recommendations

R1. For hosting providers: Implement dedicated abuse API endpoints with XARF support. Our data shows this single change correlates with an 88% reduction in TTD.

R2. For hosting providers: Ensure 24/7 abuse team coverage or automated suspension capabilities to address the 2.1× weekend gap.

R3. For abuse reporters: Adopt XARF v4 as the default report format and implement automated 24-hour re-notification with 48-hour escalation to upstream providers.

R4. For registrars: Implement rate limiting and identity verification for bulk registration APIs. Registrars with bulk APIs show 3.2× higher phishing rates.

R5. For the industry: Develop a standardized abuse reporting API specification complementing XARF, with defined endpoints for submission, status, and resolution.

R6. For policymakers: Consider regulatory frameworks requiring minimum takedown SLAs for hosting providers, analogous to the EU Digital Services Act notice-and-action provisions.

13.3 Future Work

Planned extensions include: (1) expanding to a 12-month longitudinal study; (2) analyzing the impact of reporter reputation on provider response speed; (3) developing a public provider scorecard ranking hosting providers by abuse response metrics; (4) integrating takedown telemetry into the vSpam.org DNSBL for automatic delisting upon confirmed takedown; and (5) evaluating the Digital Services Act's impact on EU-based provider response times.

References

[1] Spamhaus Technology. "Real-Time DNS Blocklists." https://www.spamhaus.com/

[2] Netcraft. "The Definitive Guide on Leading Phishing Takedown Providers." 2025. https://www.netcraft.com/blog/the-definitive-guide-on-leading-phishing-takedown-providers

[3] PhishFort. "Phishing Takedown Services." 2025.

[4] Shafranovich, Y. et al. "An Extensible Format for Email Feedback Reports." RFC 5965, IETF, August 2010.

[5] Abusix. "XARF v4: eXtended Abuse Reporting Format." https://xarf.org/

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

[7] CrowdSec. "The CrowdSec Data." https://www.crowdsec.net/our-data

[8] Interisle Consulting Group. "Phishing Landscape 2025." https://www.interisle.net/

[9] M3AAWG. "Best Common Practices for Anti-Abuse." https://www.m3aawg.org/

[10] vSpam.org. "The DNSBL Effectiveness Study." VSPAM-TR-2026-002, March 2026.

[11] Google. "Email sender guidelines." https://support.google.com/a/answer/81126

[12] CloudSEK. "5 Trusted Phishing Domain Takedown Services in 2026." https://www.cloudsek.com/

[13] phish.report. "Tools to combat brand impersonation." https://phish.report

[14] Intra2net. "Blacklist Monitor: Statistics of Accuracy and Failure Rates." https://www.intra2net.com/en/support/antispam/

Appendix A: Statistical Methods

Test / Method	Application	Parameters
Mann-Whitney U	API vs. email TTD; weekday vs. weekend	Two-sided; $\alpha = 0.05$
Kruskal-Wallis H	TTD across 5 provider tiers	H statistic; post-hoc Dunn's
Log-linear regression	TTA–TTD relationship (Fig. 11)	$R^2 = 0.71$; power-law fit
Bootstrap CI	Mean TTD confidence intervals per tier	B = 10,000; BCa method
χ^2 test	Re-notification strategy TSR comparison	χ^2; df = 3; $p < 0.001$
Wilcoxon signed-rank	Paired XARF vs. free-text at same providers	Matched-pair design

Table A1: Statistical methods applied. All tests two-sided; $\alpha = 0.05$ unless noted.

Appendix B: Nomenclature

Acronym	Full Term
API	Application Programming Interface
ARF	Abuse Reporting Format (RFC 5965)
CDN	Content Delivery Network
CERT	Computer Emergency Response Team
CDF	Cumulative Distribution Function
DGA	Domain Generation Algorithm
DNSBL	DNS-based Blocklist
FIRST	Forum of Incident Response and Security Teams
IQR	Interquartile Range
M3AAWG	Messaging, Malware and Mobile Anti-Abuse Working Group
MTA	Mail Transfer Agent
RBL	Real-time Blackhole List
SLA	Service Level Agreement
SMTP	Simple Mail Transfer Protocol
TSR	Takedown Success Rate
TTA	Time-to-Acknowledgment
TTD	Time-to-Takedown
WHOIS	Domain registration query protocol
XARF	eXtended Abuse Reporting Format

About vSpam.org

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

Document ID: VSPAM-2026-010 | Version: 1.0 | Classification: Public | DOI: 10.xxxx/vspam.2026.010 (pending)