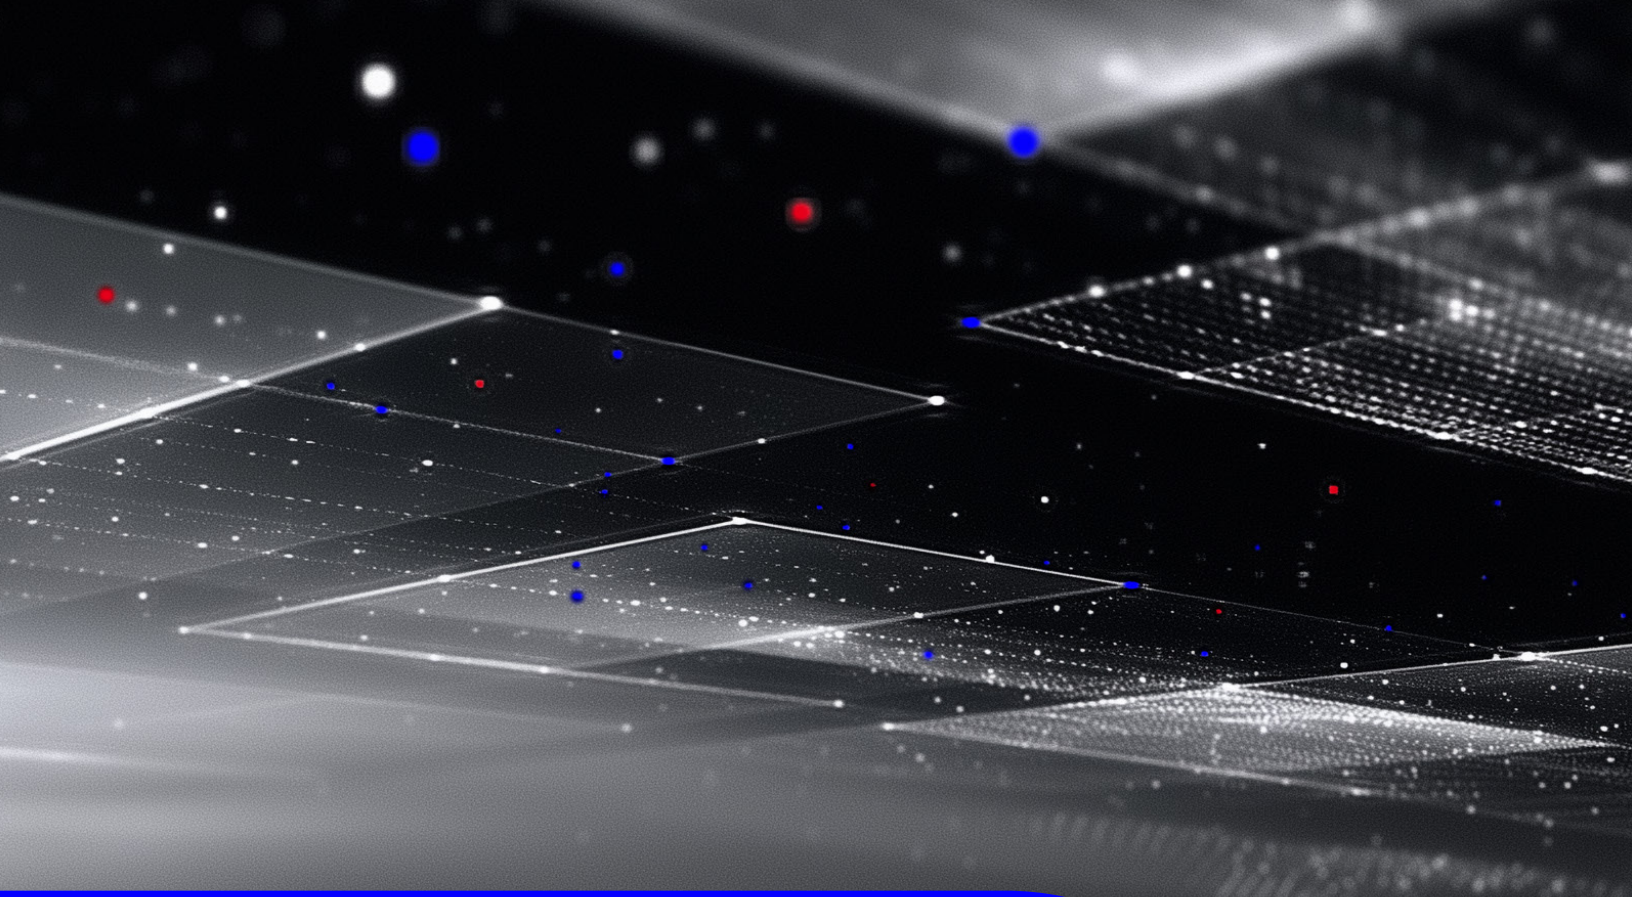

Annual Payment Fraud Intelligence Report: 2025

Recorded Future®

Expanded fraud attack surface for stolen cards and popular one-time password theft tactics raised the threat of payment fraud in 2025, especially when enabled by social engineering attacks.

Duplicate stolen check volume in the US dropped, but the count of unique stolen checks increased, suggesting that check fraud risk remains persistent for US consumers and financial institutions.

Industrial-scale support ecosystems allowed e-skimmer infections, purchase scams, and payment card testing to flourish, but growing standardization likely increases the impact of fraud intelligence.

Executive Summary

In 2025, the global payment fraud threat landscape grew more complex and dangerous due to expanding fraud attack surfaces, increasing sophistication, and the emergence of novel fraud threats. The quantity of stolen payment card data posted for sale on the dark web declined, but sensitive cardholder attributes accompanying that data increased, as did the raw volume of all freely exposed card data. Support networks for attack vectors industrialized: AI enablement and increasingly professionalized fraud-support services allowed threat actors to maximize fraud outputs while minimizing input requirements. Ready-to-go e-skimmer kits and services helped threat actors steal victim data during online transactions at scale, and purchase scam networks reached massive proportions, abusing thousands of merchant accounts as fraud attack entry points. One-time password interception cemented its popularity as a tactic to bypass strong authentication, often for digital wallet fraud. As threat actors demonstrated increasing willingness to incorporate AI enablement into their attack chains, pilot programs for agentic commerce introduced novel fraud risks and unclear fraud liability into the e-commerce industry, likely raising operational risk in 2026.

Maintaining an effective fraud posture will increasingly require complementing reactive account monitoring with proactive, intelligence-informed defenses.

Increasing AI threat enablement, the popularity of one-time password interception, and broadly expanding fraud attack surfaces across the payment and threat landscape likely amplify the risk of social engineering attacks, which in turn increase the likelihood of success for unauthorized fraudulent transactions and other downstream fraud. Ecosystem-enabled scale and sophistication for e-skimmer malware and purchase scam operations also point to a likely increase in sheer fraud volume, especially as victim-authorized fraud seen in purchase scams lacks effective corresponding controls. Furthermore, the lack of transparent fraud controls and operationally ambiguous liability in agentic commerce presents acute operational and fraud risks for merchants and financial institutions, which will likely absorb elevated fraud losses until consistently reliable agentic commerce guardrails mature.

Last year's trends suggest that going forward, maintaining an effective fraud posture will increasingly require financial institutions and payment organizations to complement reactive account monitoring with proactive, intelligence-informed defenses. Such defenses allow fraud risk stakeholders to generate a single view of linked attack behavior and risk context from signals collected across multiple domains. Achieving this will require leadership-level alignment and investment to fuse organizations' cyber threat intelligence and fraud functions. However, successful implementation — for example, through the deployment of intelligence-driven cyber-fraud fusion centers — will likely improve long-term signal identification and decision speed, especially as threat actors evolve to overcome those same improvements. Earlier-stage disruption should remain the focus of these efforts, as early stages are where intelligence enablement and anti-fraud automation can most effectively disrupt threats at scale, before financial losses come into play.

Key Findings

- **Total data exposure flatlined in 2025, but the fraud attack surface of exposed data expanded.** Threat actors posted more than 142 million stolen card records for sale on dark web marketplaces in 2025, down 19% from 2024. However, the quantity of stolen cards with accompanying contact information increased by nine percentage points to 82%, raising the risk of downstream fraud attacks, and the quantity of freely exposed payment card records on Telegram and other dark web sources increased by 26%. As the volume of stolen US paper check images posted for sale on Telegram decreased by 42% to 1.3 million, the quantity of unique stolen checks increased by 3% to 233,000, suggesting that the risk of check fraud for US financial institutions will likely remain persistent in 2026.
- **The ecosystem sustaining upstream fraud operations reached industrial proportions.** This industrialization was driven by technical advances and increasingly professionalized support services that maximize fraud outputs while reducing input requirements. Magecart e-skimmer attacks that steal data during online transactions grew more scalable and complex due to advancing e-skimmer malware kits and services, likely enabling the more than 10,500 e-skimmer infections active in 2025 to compromise over 23 million online transactions. Data patterns for more than 3,600 scam merchant accounts that were abused to defraud victims through authorized transactions indicate that scam operators have likely adopted scalable merchant account acquisition workflows. Threat actors continued to abuse major merchants for card testing prior to fraud attacks, and Telegram-based card testing services helped validate at least 27 million payment card records over the year.
- **One-time password interception cemented its popularity as a technique for circumventing authentication.** Threat actors increasingly used one-time password interception to fill gaps in technical fraud schemes, particularly when supporting digital wallet fraud and near-field communication relay attacks. The combined techniques bypass strong authentication and were supported by the expanding fraud attack surface for data along with a slew of tools and services advertised on dark web sources.
- **Threat actors enthusiastically incorporated AI enablement into fraud attack chains.** One observed purchase scam operation used an AI-powered marketing platform to improve its victim targeting capability, and threat actors on various sources commonly discussed workflows that integrated AI into their attacks. The first reported use of AI orchestration in a cyber-espionage campaign coincided with an observed fraudulent purchase attempt for the same AI platform, demonstrating how threat actors can employ fraud tactics to fund nefarious activity while protecting their identities.
- **The rise of AI-powered agentic commerce introduced novel fraud risks for the e-commerce industry.** Agentic commerce has introduced user and agent intent as a novel attack surface, which is likely vulnerable to the same attack methods that circumvent identity authentication and enable payment fraud during online transactions. It also introduces a new burden for fraud disputes: Even with clear liability, issuers face higher investigation costs, as agent behavior complicates attribution and triage.

Top Payment Fraud Risks and Strategic Implications for 2026

The 2025 payment fraud threat landscape suggests that three major themes will define fraud risk in 2026. Financial institutions will likely need to defend a wider fraud attack surface, especially as merchant accounts become fraud attack entry points for surging purchase scams and online shoppers increasingly expose their contact information alongside their financial information. As fraud ecosystems industrialize due to growing enablement and professionalization, barriers to entry for inexperienced or unskilled threat actors will likely shrink, raising the scale and sophistication of upstream attacks. Combined, these dynamics will likely generate pressure on financial institutions, merchants, and other payment organizations to incorporate intelligence-enabled fraud disruption strategies, which business leadership must strive to nourish through alignment and investment.

To counteract current threat trends in 2026, executives and business leaders should mandate cross-functional collaboration between their organizations' intelligence functions, fraud functions, and other business functions that own elements of fraud risk. Given that 2025 heralded an increase in the actionability of fraud data signals, intelligence teams should grow comfortable orchestrating information-sharing processes between fraud functions and other organizational stakeholders. To fully leverage intelligence, fraud teams should effectively incorporate intelligence processes — manually if necessary, but with automation as the desired end state — especially as threats evolve.

2026 Top Risks and Strategic Implications

	Theme 1: **Wider fraud attack surfaces**	Theme 2: **Industrializing fraud ecosystem**	Theme 3: **Broader fraud intelligence adoption**
Executive Leadership	Drive cross-functional collaboration through cultural commitment, structural organization, and investment.		
Threat Intelligence	Triage and collect relevant proactive fraud signals based on fraud partners' needs.	Leverage increasingly widespread indicators to drive intelligence collection.	Inform fraud partners of intelligence capabilities while soliciting intelligence priorities.
Fraud Operations	Develop intelligence-enabled strategies to identify at-risk accounts.	Incorporate actionable — and where able, imaginative — upstream fraud signals to drive threat disruption at scale.	Actively participate in intelligence exchange to adapt to a dynamic threat landscape.
Other Fraud Risk Stakeholders	*Adopt intelligence-forward mitigation processes that leverage the expertise of frontline fraud defenders.*		

Data Exposure

Data exposure is the fuel that sustains large-scale fraud, but not all exposed data carries equal fraud risk.

Analysis in this section is drawn from three Recorded Future® Payment Fraud Intelligence datasets:

- Compromised payment card data posted for sale on dark web marketplaces
- Full payment card data exposed on various sources, including Telegram and dark web forums
- Images of stolen US paper checks posted for sale on check fraud–focused Telegram sources

In many cases, data exposure offers threat actors an entry point into multiple fraud attacks targeting their victims. Therefore, high-quality victim data exposure not only presents an immediate fraud attack surface but also amplifies the efficiency of successive downstream fraud attacks, shortens time-to-abuse, and enables more sophisticated attack chains.

Dark Web Carding Marketplaces: Volume Declines, Attack Surface of Exposed Data Increases

Throughout 2025, fewer stolen payment card data records were posted for sale on dark web sources. Among the dark web marketplaces monitored by Recorded Future Payment Fraud Intelligence, threat actors posted over 142 million card data records for sale: 34 million fewer than in 2024, a 19% drop. Stolen payment card records support payment card fraud, often called "carding" in threat actor parlance.

At the same time, the fraud attack surface of for-sale data broadly expanded, suggesting that the lower availability of for-sale payment cards was offset by increased attack surface for those same records. Victim email addresses or phone numbers accompanied 82% of for-sale card data stolen during card-not-present (CNP) transactions, up nine percentage points from 2024. This trend points to a likely increase in attack surface for account takeovers (ATO) that incorporate social engineering techniques, and implicitly, a likely need for more nuanced controls. Even so, the trend toward attribute-rich data suggests that fraud intelligence will become more actionable over the coming year.

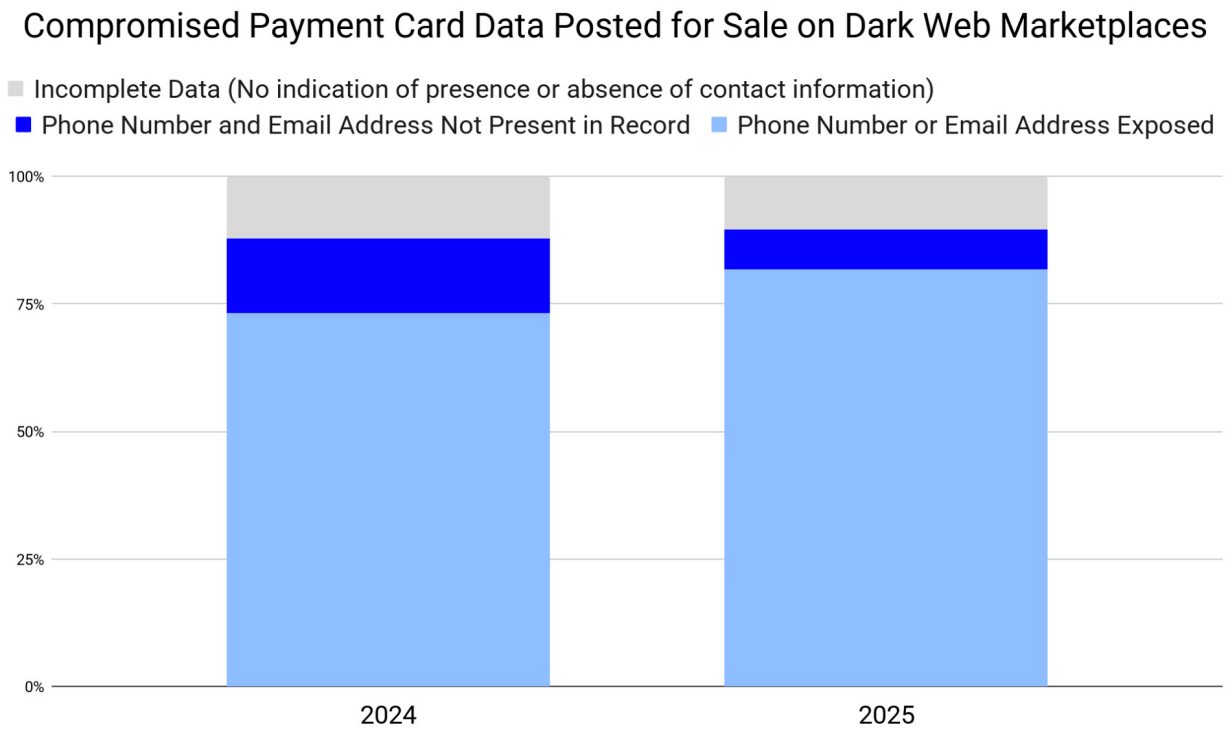

The expanding attack surface of stolen data in 2025 was most clearly evident in the increase in accompanying contact information for stolen payment card records posted for sale on the dark web.

As is typical, CNP data used for online fraudulent transactions dominated the dark web card market, but supply and attack opportunities helped card-present (CP) fraud during in-person transactions remain a potent threat. Payment Fraud Intelligence analysts documented the continuing use of point-of-sale malware "MajikPOS" in mid-2025, and a sevenfold surge in the availability of stolen Polish CP data accompanied an October spike in ATM fraud. Altogether, more than 21 million stolen CP card records were posted for sale in 2025, 94% of which were US-issued.

Despite Market Disruptions, the Dark Web Carding Economy Perseveres, and Total Volume of Free Card Data Increases

Market disruptions were partially responsible for the reduced volume of card data offered for sale on the dark web in 2025. Law enforcement intervention led to the closure of the dark web marketplace "BidenCash" and interruptions for multiple smaller marketplaces. Before closing, BidenCash continued its tradition of releasing large databases of free card information, a marketing tactic that the source "b1ack's Stash" mimicked. The seizure of BidenCash followed the marketplace's steady decline in market share. Other takedowns affecting at least five card data marketplaces had less impact on the greater threat landscape: Most of the shops continued to post stolen card data records on related domains, demonstrating resilience as fraud epicenters.

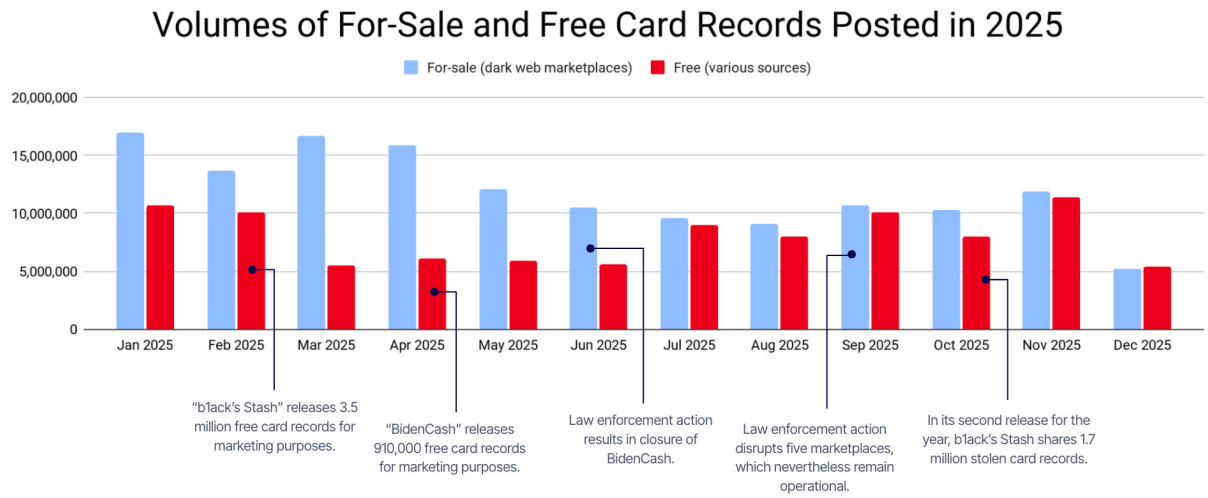

As for-sale card volume on the dark web decreased, the average monthly volume of freely exposed records increased compared to 2024. Threat actors freely exposed more than 142 million card records on various sources, up 26% from 2024 and comparable in volume to the total quantity of stolen payment cards posted for sale on the dark web. Although many of these freely exposed card records lacked the accompanying personally identifiable information that enables threat actors to conduct CNP fraud, such as name and address, the inclusion of contact information with many free card records nevertheless likely increased the risk of spearphishing attacks and online banking ATO attacks.

Regional analysis indicated market stability despite pressure. The volume of stolen payment cards available for sale on dark web marketplaces declined in nearly every region compared to 2024, although to a lesser extent in the Middle East, Türkiye, and Africa. Europe was the major outlier for the period. In 2025, the volume of stolen European payment cards posted for sale on the dark web increased by nearly 1.5 million records compared to 2024.

Top Countries per Region by Dark Web Payment Card Exposure, 2025

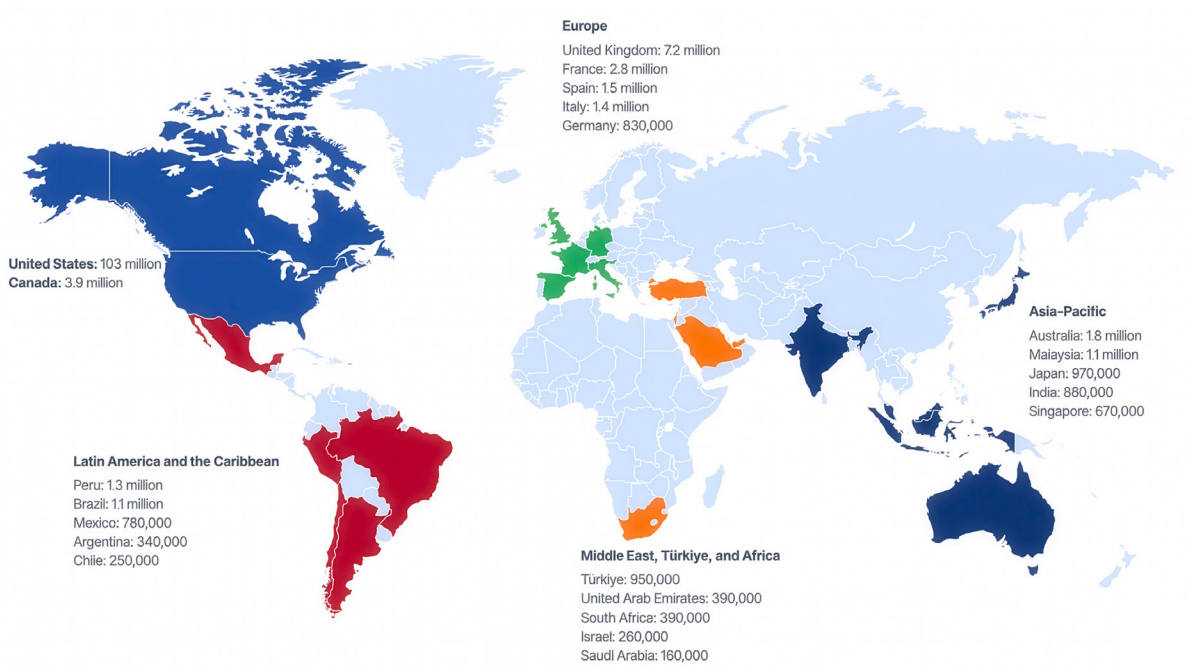

Total Volume of Stolen US Bank Checks Decreases and Deurbanizes, Signaling Potential Shift in Mail Theft Trends

The volume of stolen US payment check images exposed on Telegram decreased by 42% in 2025. Although the decrease follows Telegram's late 2024 agreement to share user data with authorities, analysis from Recorded Future attributed most of the decrease to targeted action by Telegram against channels that violated policy. Typically, paper checks are stolen from mail collection points[1] before images of them are posted for sale on Telegram.

1.3 million
total US payment check images posted for sale in 2025

42%
fewer check images posted for sale, a decrease entirely attributable to fewer reposted, duplicate images

233,000
unique US payment check images posted for sale for the first time, up from 2024

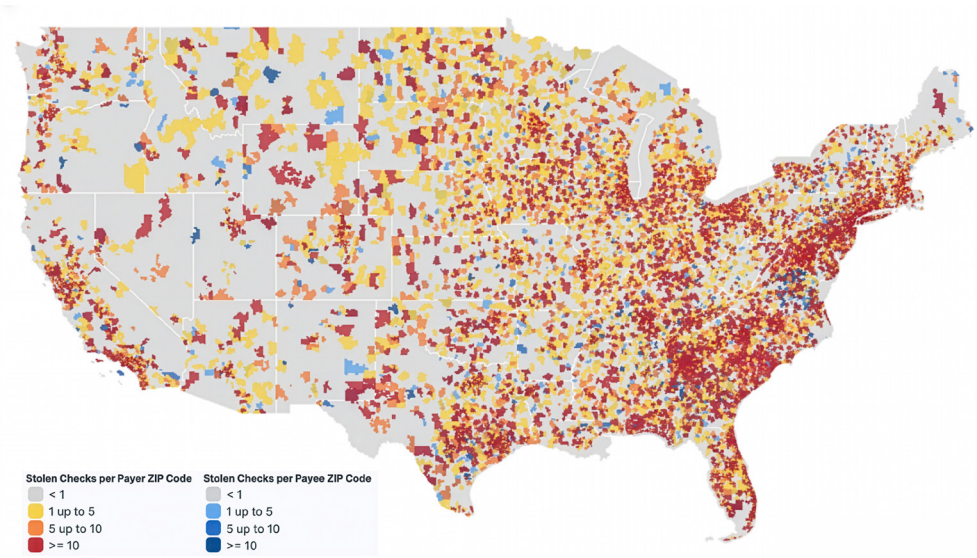

As fewer paper check images were posted for sale in 2025, the geographic distribution of those stolen checks became less focused around major urban centers.

The drop in total check image volume is unlikely to herald a corresponding drop in check fraud. Stolen check images are duplicated and reposted across multiple sources following their first posting. In 2025, the observed drop was entirely attributable to fewer duplicate check images. The volume of unique check images posted for sale actually increased by 3% over 2025, indicating that policy changes and targeted action likely had little impact on the activity of stolen check suppliers. As fewer check images were reposted, trends in geographical origin shifted, suggesting that criminal groups engaged in mail theft have adjusted their mail theft strategies. The volume of checks stolen from payers outside of major urban centers increased compared to 2024, particularly in the US Midwest.

1 In January 2025, the US Federal Bureau of Investigation and US Postal Inspection Service (USPIS) released a joint alert regarding mail theft-related check fraud. The alert follows a similar alert released in February 2023 by the Financial Crimes Enforcement Network and USPIS.

The possible shift in mail theft trends, combined with the stable volume of new stolen checks, indicates that the US check fraud threat landscape is undergoing a transition that could potentially impact downstream check fraud patterns and controls. While refinement of controls will likely remain useful in 2026, the implementation of automated, intelligence-based mitigations using structured check data will likely be more effective in helping financial institutions weather the change.

Attack Vectors

Fraud outcomes are visible, but the pathways that enable them are often not.

This section examines observed trends in attack vectors used to target merchants and consumers in 2025, as well as what those trends likely mean for payment fraud risk going forward. As a rule, common attack vectors reveal where attacker behavior is converging. By mapping these vectors to downstream fraud outcomes, anti-fraud practitioners can understand how to leverage early-stage attack indicators to disrupt threats before impact manifests financially.

Analysis in this section is drawn from three Recorded Future Payment Fraud Intelligence datasets:

- Merchants with websites compromised by Magecart e-skimmers that steal customer data during online transactions
- Tester merchants abused to test stolen payment cards before the execution of the final fraud attack
- Scam merchants linked to scam websites that defraud victims through authorized transactions and data theft

Understanding attack vectors helps clarify control and visibility gaps early in the fraud lifecycle, before financial losses are recorded. These findings can help business leaders prioritize defensive investment, fraud teams refine detection, and cyber threat intelligence teams identify and address emerging techniques before they scale.

Upstream Fraud Attack Methodology Advances as the Broader Fraud Ecosystem Industrializes

Recorded Future Payment Fraud Intelligence monitors e-commerce websites for indicators of "Magecart" e-skimmer infections: malicious scripts that steal victim data during online transactions. Magecart e-skimmer attacks occur upstream of fraud attacks, so early identification of e-skimmers enables fraud teams to proactively detect customer data compromise.

The most prominent theme in the 2025 Magecart threat landscape was the advancement of full-stack e-skimmer kits and Malware-as-a-Service (MaaS) offerings into a more industrialized enablement ecosystem. These toolsets lowered barriers to entry for Magecart threat actors and reflected a broader shift toward scalable fraud enablement that likely accelerated e-skimming operations and compressed the time between victim data compromise and monetization throughout the year.

Overall, Magecart impact did not appear to decrease following the implementation of PCI DSS 4.0's future-dated requirements in early 2025. While total persistent and new infection

volume dropped, the decrease was relative to a massive 2024 spike in infection volume that resulted from CVE-2024-34102, a vulnerability also known as "CosmicSting." Altogether, Magecart e-skimmer infections active in 2025 likely compromised more than 23 million transactions.[2]

10,500+
total unique Magecart e-skimmer infections active at any point in 2025, stable compared to 2024

7,300+
new unique Magecart e-skimmer infections detected at any point in 2025

23.4 million
transactions likely compromised by Magecart e-skimmer infections in 2025

Among all e-skimmer malware packages offered for sale on the dark web in 2025, "Sniffer by Fleras" (also identified as "Surki") saw the widest usage, with 26% of all e-skimmer infections attributable to the kit.[3] Sniffer by Fleras includes access to a web-based portal capable of generating malicious scripts and a server for managing stolen data. Threat actors on dark web sources broadly consider the e-skimmer kit one of the most popular on the market.

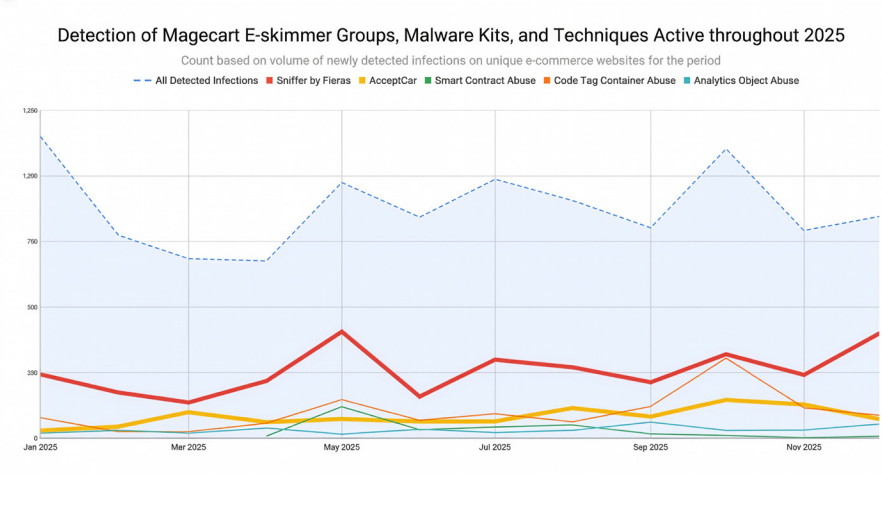

The "Sniffer by Fleras" e-skimmer kit was widely used across all observed Magecart e-skimmer infections in 2025, lowering barriers to entry for Magecart threat actors. Other notable infection patterns included the use of Magecart MaaS offerings — as seen with "AcceptCar" — and abuse of various legitimate services.

The "AcceptCar" e-skimmer was discovered in H2 2025, with the number of infected websites increasing towards the end of the year. The e-skimmer is operated under a Magecart-focused MaaS model: In return for e-skimmer installation and operation on compromised e-commerce websites, threat actors reimburse AcceptCar operators with 50% of proceeds from sales of stolen card data or 70% of raw data intake. AcceptCar operators protect their e-skimming scripts with robust obfuscation methods, which complicates the analysis and extraction of attack indicators.

2 This estimate is a minimum and is based on analysis of infection volume, aggregate infection periods, and average monthly visitor statistics.
3 Recorded Future Payment Fraud Intelligence analysts first identified the e-skimmer kit in H1 2024

Historical Magecart attack patterns remained relevant in 2025. Magecart threat actors continued to exploit legitimate services for their e-skimmer attack chain, using code-tagging containers for loader, relay, and e-skimmer scripts. Newly identified techniques involved the abuse of embedded smart contract scripts to house e-skimmer payloads and mimicry of analytics objects to avoid detection. Magecart breaches continued to plague centralized e-commerce platforms. One case impacted a reservation platform, affecting over 160 business customers after e-skimmers were injected into client-specific code repositories.

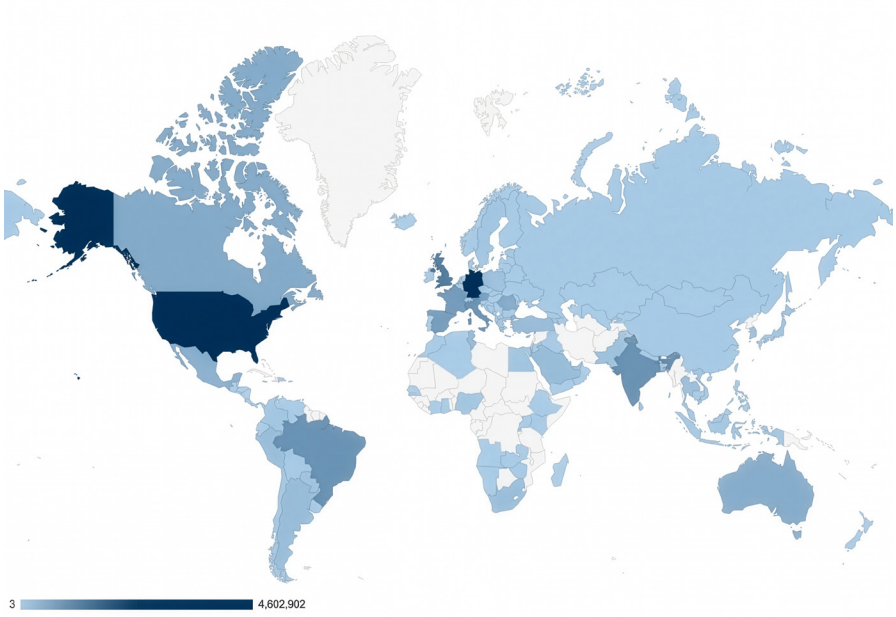

Magecart e-skimmers likely compromised more than 16 million transactions in 2025. The map chart above illustrates compromise estimates, with darker colors signifying a higher concentration of data compromise.

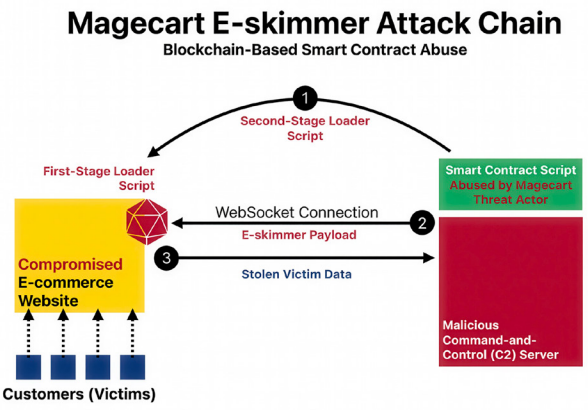

After gaining illicit access to e-commerce websites, a Magecart threat group implanted first-stage loader scripts that retrieved a second-stage loader from a blockchain-based smart contract, which in turn opened a WebSocket connection to retrieve the e-skimmer payload. Stolen data was exfiltrated via the same WebSocket connection.

The increasing use of Magecart e-skimmer kits and services in 2025 suggests that in 2026, attacks will become more widespread and more complex. Moreover, stable infection volume following the implementation of PCI DSS 4.0's future-dated requirements indicates that compliance rules alone are likely insufficient to remediate the industry-wide impact of e-skimmer infections. Although these factors point to amplified downstream fraud risks for the attack vector, the growing standardization of the e-skimmer supply chain also indicates that proactive detection of Magecart indicators will likely remain effective, even as increasing complexity makes the discovery of those indicators more difficult.

Sophisticated Purchase Scam Website Networks Tap Fraudulent Merchant Accounts, Achieving Massive Scale

Online purchase scams enabled by scam websites emerged as a serious threat in 2025.[4] Simultaneously, Payment Fraud Intelligence detections of purchase scam merchant accounts more than quadrupled compared to 2024, demonstrating once again how payment infrastructure can be subverted to achieve fraud outcomes. In a purchase scam, visitors attempt to buy non-existent goods or services from fake e-commerce websites, authorizing payments to fraudulent merchant accounts in the process. The rapid emergence of purchase scams as a threat is likely a threat actor response to improving authentication controls that increase resource costs for unauthorized fraudulent CNP transactions. Purchase scams neatly circumvent these controls by manipulating cardholders into authorizing the fraudulent transactions themselves.

3,600+
scam merchant accounts identified in 2025, up 2.5x from 2024

40+
countries associated with all identified scam merchants in 2025

230+
acquirers associated with all identified scam merchants in 2025

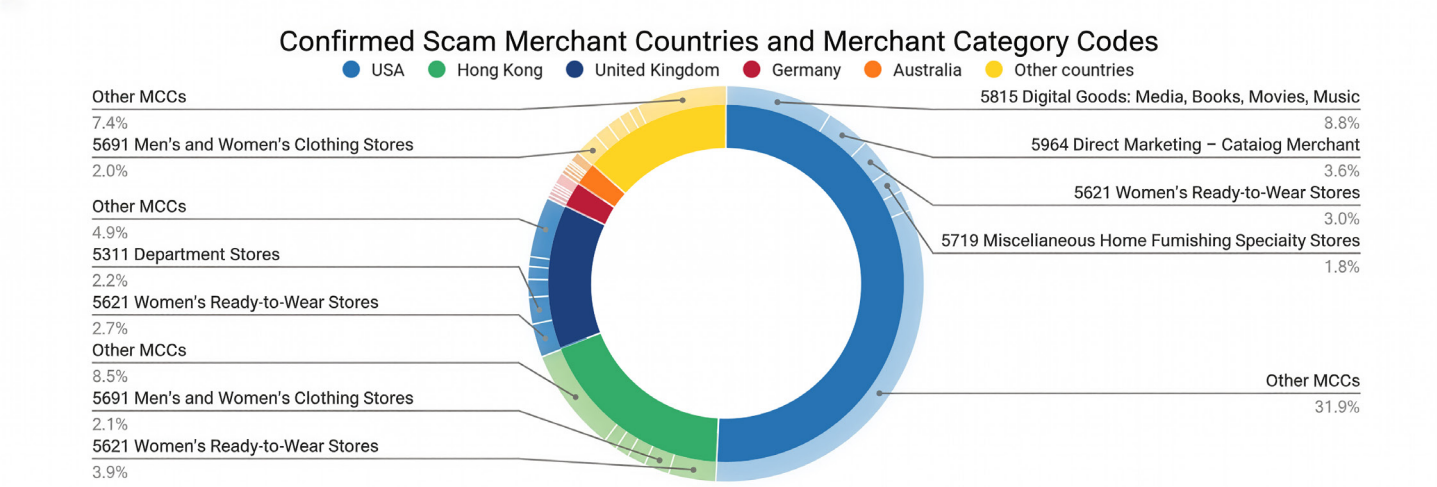

Recurring patterns in merchant registration data suggest that threat actors devise and employ reliable account acquisition workflows to stand up their merchant infrastructure.

Analysis this year indicates that active purchase scam ecosystems have achieved massive scale and sophistication. Data compromise, a hallmark of purchase scams, was confirmed by links between more than 50 identified scam merchants and for-sale card data on the dark web. The largest scam campaigns tricked victims into enrolling in subscription traps, where their

4 Mastercard released a scam-focused whitepaper in May 2025, and joint research released by the Global Anti-Scam Alliance and Feedzai in 2025 indicated more than $1 trillion in all scam losses throughout 2024.

payment cards are billed on a recurring basis. An additional tactic observed in H2 2025 involves immediate secondary charges against a victim's payment card from a purported transaction recovery service, likely to bypass fraud controls. Common patterns between acquirer, country, and scam merchant data, such as merchant category code, suggest that threat actors likely leverage scalable workflows to acquire their fraudulent scam merchant infrastructure.

Analysis of scam domains throughout the year indicated consistent brand and ad abuse for identified scam operations. Social media advertisements are the primary targeting mechanism used to attract victims to scam websites:[5] Threat actors pay for ads on social media platforms, whose marketing algorithms target users likely to be interested in purchasing the scam's purported offering. To increase the success of their advertisements, purchase scam operators impersonate popular brands and advertise steep discounts.

2025 threat trends suggest that agentic tooling will likely catalyze a new evolution in social engineering throughout 2026, lowering skill barriers to establish and operate purchase scam campaigns. One scam operation observed in 2025 incorporated an AI-powered marketing platform to streamline phishing message generation and victim targeting. On November 29, 2025, a detailed tutorial on fraud-focused Telegram channels offered guidance for building fully functional phishing pages using autonomous AI agents. Beyond web templates, threat actors increasingly discussed using generative AI systems for voice-cloned vishing campaigns, AI-powered call-center scripts, and deepfake identity artifacts to support ATO and real-time social engineering attacks.

Purchase scams' ability to circumvent traditional CNP fraud controls indicates that scam detection will likely remain an enduring challenge — and should become a strategic imperative — for financial institutions in 2026. Solutions must not rely solely on traditional fraud detection models, which threat actors can defeat by eliciting victim authorization and pivoting to new infrastructure. Predictive and intelligence-based strategies that proactively identify merchant relationships with scam networks will likely be more reliably effective, especially as threat actors continue scaling their scam activity.

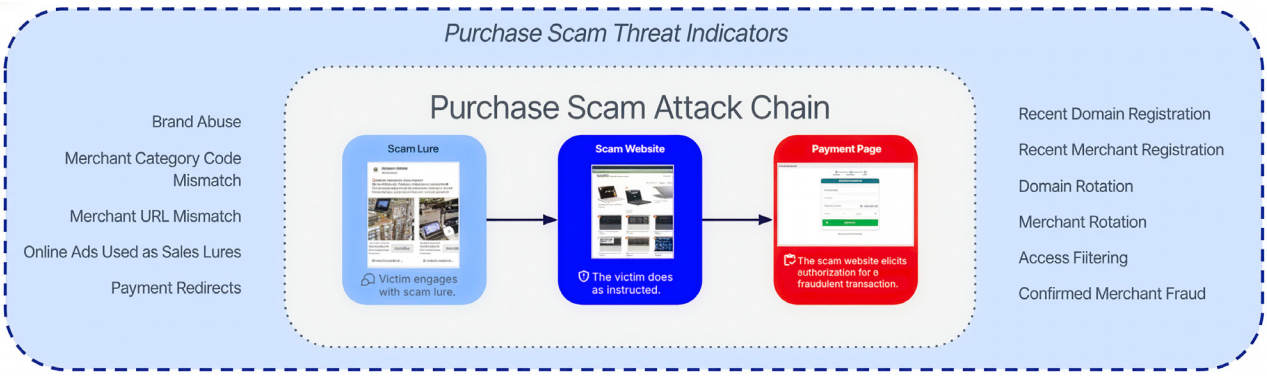

5 The purchase scam ecosystem is deeply intertwined with the greater online ad ecosystem, as reported by Reuters in November 2025.

One-Time Password Interception Increasingly Fills in Gaps for Fraud Workflows That Abuse Payment Technology

Threat actors increasingly used one-time password (OTP) interception to fill gaps in technical fraud schemes over the past year. This is part of a continuing trend reported by multiple sources in recent years, including Recorded Future in 2024 and UK Finance in May 2025. Similar to purchase scams, OTP interception's growing use is likely a threat actor response to the challenge of bypassing strong authentication driven by the broad use of OTPs for secondary verification.[6] Expanding attack surface for exposed data in 2025 likely contributed to the popularity of OTP theft as a fraud tactic.

Digital wallet fraud often requires OTP interception for success. Threat actors load stolen card data onto digital wallets, then conduct "trusted" fraudulent transactions using the wallet. Near-field communication (NFC) relay attacks are an extension of digital wallet fraud that gained prominence in 2025: Multiple sources reported surges in "ghost-tapping" attacks, during which threat actors relay stolen card data to money mules' devices remotely to support unauthorized contactless transactions. As wallet-based fraud schemes matured throughout the year, threat actors on dark web sources advertised more and more sophisticated OTP interception tools and services to support them.

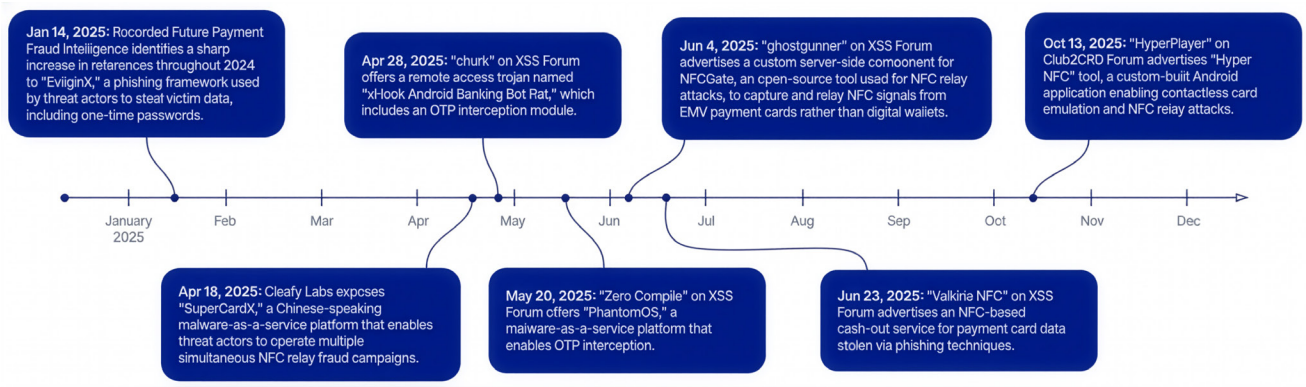

A slew of dark web offerings supported digital wallet fraud and NFC relay attacks in 2025.

The impact of OTP interception's growing popularity and the vital role it plays in other attack chains was evident throughout the year. In June 2025 Recorded Future Payment Fraud Intelligence analysts investigated a high-volume phishing campaign that targeted UK residents with phishing lures modeled after winter fuel payment notifications. Attacks culminated in OTP theft attempts that supported digital wallet fraud, and the campaign was linked to stolen card data. Meanwhile, the Magecart threat group "OTPExplorer," first identified in 2024, continued deploying e-skimmers that incorporate OTP theft functionality, likely to facilitate digital wallet fraud.

Since wallet transactions are often trusted, the growing momentum of wallet-based fraud and NFC relay attack methods in 2025 will likely complicate fraud detection and increase CP fraud risks throughout 2026. Although the OTP interception techniques that enable these attacks

6 A 2025 TransUnion report identified OTPs as the most common form of secondary authentication globally.

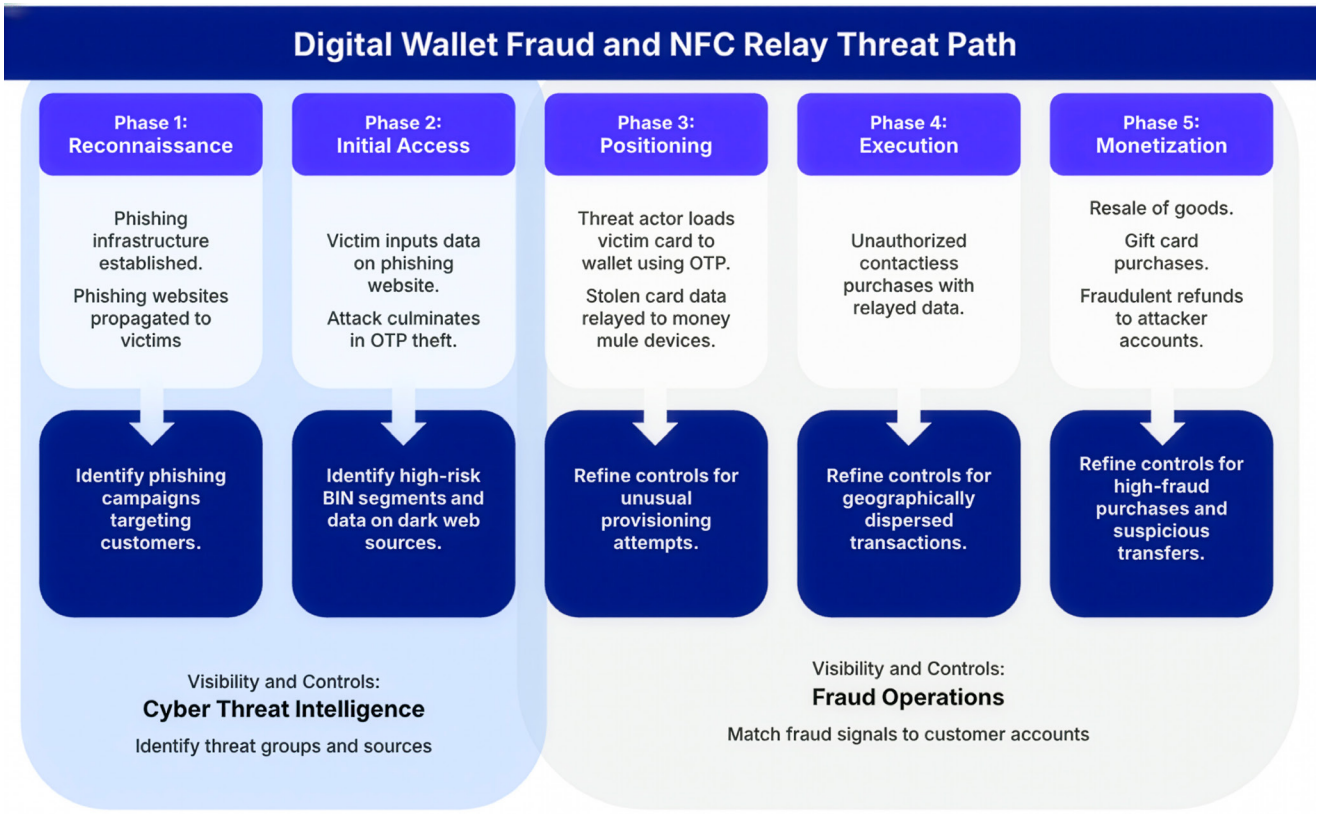

Close collaboration between intelligence and fraud operations teams can disrupt OTP interception–based attack methods.

ultimately amount to social engineering, human error is inevitable. Consequently, the natural OTP interception attack surface will likely remain difficult to minimize. Fraud defenders in 2026 should look to the more reliably detectable disruption point for these attack vectors: the card provisioning attempt, where refined fraud controls can detect suspicious provisioning requests, especially when buoyed by proactive intelligence signals.

Dark Web Threat Actors Increase Their Access to Card Testing Infrastructure through Tester Merchant Abuse

Threat actors on the dark web actively increased their access to card-testing infrastructure, compromising or abusing more than 1,350 tester merchants throughout 2025. As a rule, threat actors

1,350+
tester merchant accounts identified in card testing activity

94%
of tester merchant accounts identified in 2025 were not previously observed

27 million
card data records on Telegram were exposed during card validation attempts

preferred new tester merchants: 94% of all tester merchants identified in 2025 were not observed previously, up four percentage points from 2024. The abuse of new tester merchants enables threat actors to evade detection, especially as historical tester merchants are detected and flagged for fraud.

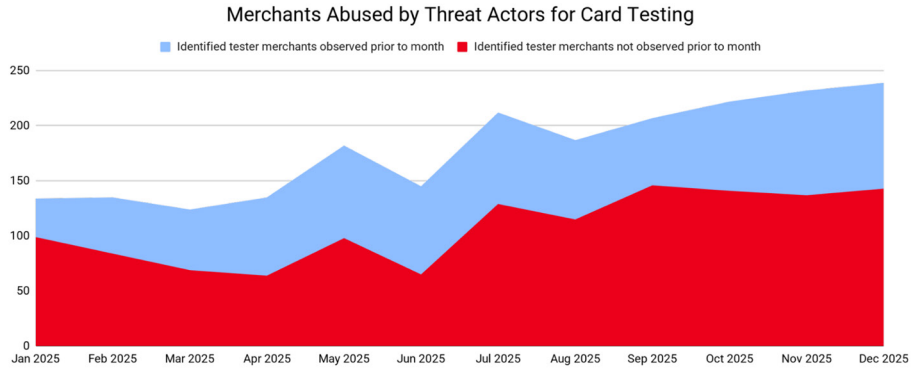

On a month-by-month basis, most tester merchants observed were new; 94% of all tester merchants observed in 2025 were not observed in previous years.

Telegram data also indicated a high volume of card testing activity in 2025. More than 27 million card records — 38% of all cards observed by Recorded Future on Telegram sources — were exposed on Telegram channels that offer public-facing card generation and testing services, which can support bank identification number (BIN) attacks. BIN attacks allow threat actors to generate valid card numbers using specific BINs for downstream fraudulent transactions.

Abuse of major legitimate merchants for card testing continued, with threat actors abusing major retailers' payment infrastructure to validate stolen cards. Recorded Future Payment Fraud Intelligence reported on seven tester services joining this trend in 2025. These services abused merchants that ranged from an e-commerce platform serving schools to a US state court payment system.

Blanket fraud controls are often ineffective for large merchants that are abused for testing. Therefore, tester merchant intelligence will likely remain effective for card testing, detection, and remediation strategy in 2026.

Agentic Commerce Poses New Fraud Risks and Liability Ambiguity for Financial Institutions

The pressure to validate security across multiplying agentic commerce services will likely define operational payment friction in 2026. This challenge emerged in 2025 as the payments industry piloted agentic commerce systems, which allow AI agents to make purchases on consumers' behalf. Amazon's Buy for Me, Visa's Intelligent Commerce, and Mastercard's Agent Pay marked these early efforts.

The crucial complication that stems from agentic commerce is the added element of user and agent intent, which can be considered analogous to identity. Intent authentication frameworks exist and will likely experience rapid improvement to buttress weaknesses going forward, but

intent nevertheless creates a novel fraud attack surface that is almost certainly susceptible to abuse. Intent authentication will likely remain vulnerable to circumvention techniques that parallel the techniques used to defeat identity verification processes.

Control visibility gaps reminiscent of the early open banking era compound authentication challenges. While liability for agentic fraud may be contractually defined, this does not reduce the new operational burden introduced by fraud disputes:

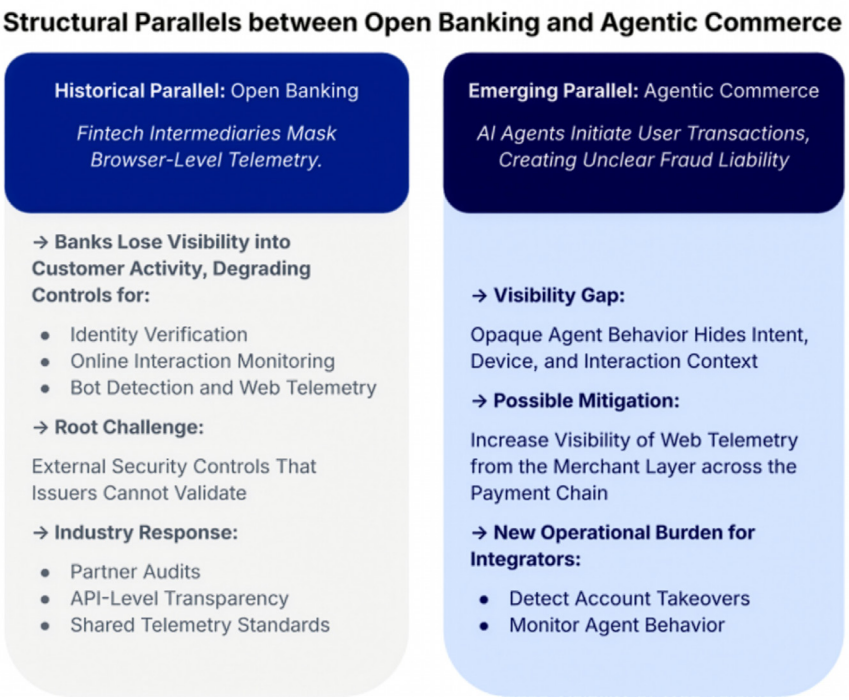

- Agentic transactions likely expand investigation scope, requiring additional resources to distinguish third-party fraud, first-party misuse, and agentic abuse.
- Clear liability does not eliminate issuer costs if attribution is contested and agent behavior must now be assessed.
- Without improved control visibility and shared investigation workflows, ecosystem players not directly liable for agentic fraud will likely resist absorbing these operational costs.

Over the past year, threat actor experimentation with agentic abuse moved from theory to practice. In April 2025, the threat actor "d0ctrine" published proof-of-concept documentation that outlined how agentic commerce creates openings for agent-automated fraud workflows, mimicry of human interaction patterns, and evasion of behavioral or device-based fraud controls. The post prompted broader discussions regarding agentic fraud. By November, the threat manifested operationally: Anthropic disclosed the first known cyber-espionage campaign orchestrated primarily by an autonomous AI system, which coincided with an attempted fraudulent purchase observed for the same AI service by Payment Fraud Intelligence analysts. The two events demonstrate the convergence of fraud and cyber threats: Fraud tactics offer threat groups elegant mechanisms to fund illicit activity while shielding their identities.

Agentic commerce is in a formative phase that will require careful oversight. In 2026, this uncertainty is poised to present acute operational and fraud risks. Various stakeholders will likely need to absorb agentic commerce fraud losses without clear liability guidance as detection technology is recalibrated to distinguish legitimate and malicious agent intent. Until these guardrails mature, agentic commerce will likely pose the same core challenge seen during the early open-banking era: elevated fraud exposure compounded by ambiguous financial liability.

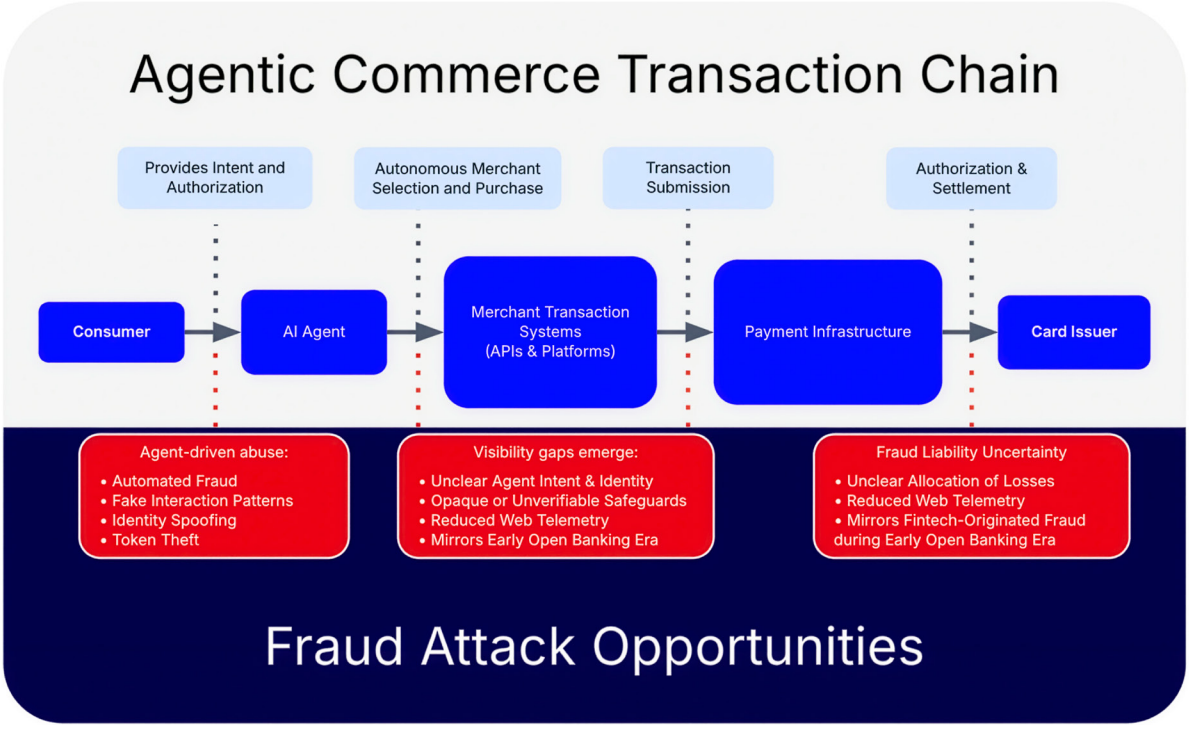

The typical agentic commerce workflow is vulnerable to fraud abuse at various points.

Mitigations

This section organizes recommended fraud mitigations around priorities for business leadership, threat intelligence, and frontline fraud operations. Effective fraud mitigation requires alignment across these three audience layers in a continuous cycle: leadership sets risk tolerance and direction, cyber threat intelligence identifies and explains emerging threats, and fraud teams operationalize controls to reduce losses. The recommendations are intended to be actionable at each level and reinforce one another across the organization.

Executives and Business Leaders

For executives and business leaders, the priority is synergetic organization, risk ownership, and investment alignment. This will require multiple initiatives: mindful structuring of operational processes that cover gaps in anti-fraud visibility, a business-as-usual mindset of proactive fraud risk management across the enterprise, and investments aimed at fortifying areas of high fraud exposure and strategic fraud risk.

- **Use Recorded Future Payment Fraud Intelligence.** Payment Fraud Intelligence offers datasets of early fraud signals. Financial institutions and other organizations can use these signals to proactively identify and remediate heightened fraud risks for customer accounts. Payment Fraud Intelligence reporting provides insights into threat actors and tactics, how and why they pose fraud risks to organizations, and recommended mitigations for addressing them.
- **Institutionalize intelligence-driven fraud operations.** Threat intelligence can enable automated anti-fraud action, but the two functions lack natural alignment in traditional business structures. Implement clear accommodations and directives to establish and sustain cooperation between intelligence and anti-fraud assets. Where necessary, leverage investment and reorganization to improve collaboration. The ultimate goal is to overcome the Castle Dilemma — the organizational difficulties that make CTI–fraud fusion challenging — by generating a single view of linked attack behavior and risk context from signals collected across multiple domains.

Cyber Threat Intelligence Teams

For cyber threat intelligence teams, the priority is early visibility, context, and effective communication with fraud partners. Fraud intelligence is highly actionable for addressing the fraud threat trends identified in this report. Cyber threat intelligence assets must therefore act as facilitators between all stakeholders who benefit from fraud intelligence within an organization.

- **Use Recorded Future Payment Fraud Intelligence.** Payment Fraud Intelligence reporting allows Recorded Future customers to identify threat groups, sources, emerging tactics, and data exposure that present fraud risks to their organization. Recorded Future customers can triage and collect proactive fraud signals from Payment Fraud Intelligence datasets, which can be packaged and delivered to fraud operations partners.

- **Establish dialogue with anti-fraud partners to prioritize actionable intelligence.** To effectively remediate at-risk customer accounts, anti-fraud assets require highly specific data signals that can be correlated to their organization's internal customer portfolio with a high degree of confidence. These data signals can vary based on the fraud item, the attack surface, and the tactics or exposed data in play. Establish a recurring dialogue with anti-fraud partners to continually apprise them of intelligence capabilities and solicit intelligence priorities that can enable proactive fraud mitigation.
- **Integrate wider stakeholders across the organization.** While tactical fraud detection and remediation occurs at the fraud operations level, other business assets — such as fraud strategy, customer service, corporate investigation, and identity and access management teams — can directly benefit from intelligence deliverables. Given their capabilities and responsibilities, cyber threat intelligence teams are well-positioned to serve as the natural connecting point between these assets.

Fraud Teams

For fraud teams, the priority is swift disruption of fraud activity at scale, particularly through intelligence-enabled automation. Given that threats will continue to evolve, intelligence enables fraud teams to routinely maintain "decision advantage" against threat actors. However, fraud teams must first effectively incorporate intelligence processes and enthusiastically collaborate with peer intelligence teams.

- **Use Recorded Future Payment Fraud Intelligence.** Payment Fraud Intelligence datasets include payment fraud signals that can be combined with automation to identify at-risk customer accounts within an organization's customer portfolio at scale, before financial losses come into play.
- **Establish intelligence-driven defense strategies.** To move from a reliance on reactive transaction monitoring and fraud models to proactive detection and prevention, incorporate the principles of threat intelligence into fraud detection and mitigation strategy at a fundamental level.
- **Adopt intelligence-forward culture and processes.** Fraud and cyber threat intelligence teams must operate as equal partners — actively sharing data, exchanging feedback, and jointly shaping intelligence requirements — so that intelligence functions as a force multiplier rather than a procedural burden. This cooperation is essential to sustain a competitive edge, especially as the threat landscape evolves.

Outlook

Looking forward, a number of specific trends will likely define the future fraud threat landscape. Despite reduced overall payment card exposure on dark web marketplaces, expanding fraud attack surfaces and the growing use of OTP interception and digital wallet–enabled fraud indicate a likely increase in CNP fraud enabled by social engineering attacks. Paper check fraud will likely remain an enduring threat in the US, regardless of whether duplicate check image volume continues to contract in the coming year. Magecart e-skimmer attacks targeting CNP payment card data for compromise will likely become more widespread and more complex as a result of the increased availability of MaaS and packaged e-skimmer malware offerings. Similarly, purchase scam operations and the victim-authorized fraud they enable, both of which have already reached massive scale, will likely persist as an evergreen fraud threat. And finally, the rise of AI-powered agentic commerce in 2025 will likely introduce novel fraud attack surfaces into the payment ecosystem, ultimately scaling the operational costs necessary to investigate and resolve agentic fraud disputes.

Stated plainly, emerging tactics and technologies, growing sophistication and expanding attack surfaces, and industrializing fraud ecosystems will likely make data exposure more actionable for threat actors and more dangerous for fraud defenders in 2026.

Nevertheless, industry defenses will also likely improve in 2026. A June 2025 Mastercard report indicated that AI-powered defenses save millions of dollars if employed well. Predictive solutions will likely become more essential as threat actors continue to incorporate AI-powered toolsets into their attack methodology. Alongside AI defenses, intelligence adoption will likely be a major theme of the coming year, especially as the expanding attack surface of exposed data makes fraud signals derived from that data more actionable for financial institutions. The growing scalability of attack vectors will also likely amplify the impact of intelligence signals: Growing standardization among Magecart e-skimmer infections due to MaaS offerings and e-skimmer kits means that indicators of compromise will surface more infections, and AI-enabled scale for purchase scam operations could simplify the detection of those operations by threat intelligence teams.

Recorded Future reporting contains expressions of likelihood or probability consistent with US Intelligence Community Directive (ICD) 203: Analytic Standards (published January 2, 2015). Recorded Future reporting also uses confidence level standards employed by the US Intelligence Community to assess the quality and quantity of the source information supporting our analytic judgments.

Recorded Future is the world's largest intelligence company. The Recorded Future Intelligence Operations Platform provides the most complete coverage across adversaries, infrastructure, and targets. By combining precise, AI-driven analytics with the Intelligence Graph® populated by specialized threat data, Recorded Future enables cyber teams to see the complete picture, act with confidence, and get ahead of threats that matter before they impact your business. Headquartered in Boston with offices around the world, Recorded Future works with more than 1,900 businesses and government organizations across 80 countries.

Learn more at **recordedfuture.com**.