When malware gets smarter and testing becomes essential

ATMs in the Age of AI

In colaboration with Cybersecurity Partner:



Table of content

03. Intro:...and (quality) testing becomes essential

04. The threat landscape: what the data says

06. The evolution of malware and the catalytic role of AI

08. The weakest link: untested software in production

12. Pre-production testing as a strategic response

- What effective testing should cover
- Serquo + Vigilis Defense: combined coverage

13. Framework of recommendation

14. Vigilis Defense and Serquo partnership

... and *(quality)* testing becomes essential

ATMs have evolved from routine financial infrastructure into one of the most active attack vectors in the global financial ecosystem

The convergence of three trends is redefining the sector's risk profile:

- The sustained rise in physical attacks and ATM fraud.
- The evolution of specialized malware toward AI capabilities.
- The democratization of offensive tools thanks to generative AI.

In this context, comprehensive software testing prior to production is no longer merely a best practice but has become a critical and indispensable security control.

Key data 2025/26

57317 incidents at CCMIS ATMIA Q1 2026

99% of ATM crime in Q1 2026 was fraud*

+89% AI-powered attacks in 2025 (CrowdStrike)



x 4.5

"Higher success rates for social engineering attacks using AI compared to traditional methods" (CertiK)

*Europe accounts for the vast majority of incidents (98%). The regional breakdown varies significantly.



The threat landscape: what the Data says

The FBI documented more than 700 jackpotting attacks in 2025 alone, resulting in losses exceeding $20 million in the U.S. In Europe, EAST recorded a 106% increase in terminal incidents between 2023 and 2024, although losses fell from 173 million to 71 million euros, a sign that defenses are improving but attackers are adapting.

ATMIA's CCMIS system, published in April 2026, accumulates 57,317 unique incidents. In Q1 2026, 99% of global ATM crime was fraud. That figure is dominated by Europe, which accounts for 98% of the quarter's incidents (1,957 of 1,992) and records 100% as fraud. North America, contributing just 20 incidents, shows the opposite profile: 90% physical attacks. Both figures are accurate and reflect dramatically different regional realities.

Critical regional divergence

The Q1 2026 CCMIS data reveals radically different risk profiles by region, with direct implications for testing strategies:

Regional differences Q1 2026

🇪🇺 **Europe:**
100% fraud, 0% physical. Cash trap 54%, card theft 45%.

🇺🇸 **North America:**
90% physical attacks. Jackpotting accounts for 50% of fraud

🌏 **Asia:**
100% card theft via card swapping.

Types of fraud Q1 2026 vs. Q4 2025

Type of attack	Q1 2026	Q4 2025
Cash Trapping	53,6%	69,0%
Card theft	45,5%	↑ 7,6%
Skimming / Data theft	0,7%	22,9%
Jackpotting	0,1%	0,5%

Source: ATMIA CCMIS Q1 2026, April 2026



> Fraud accounts for 99% of global ATM crime in Q1 2026, driven almost entirely by Europe.
>
> North America shows the mirror image: 90% physical attacks, with jackpotting leading local fraud.

Source: ATMIA CCMIS Q1 2026, April 2026

The evolution of malware and the catalytic role of AI

3 Generations in 12 years

ATM malware has undergone a generational evolution that perfectly illustrates how attackers are industrializing their methods:

Gen	Period	Features
1º	2013-2018	Manual, vendor-specific, physical activation
2º	2018-2023	Multi-vendor, Windows as a vector, remote activation
3º	2024-now	Polymorphic AI, MaaS, self-modifying code

Ploutus, detected in Mexico in 2013, required a constant physical presence. Ploutus-D (2018) was already operating on devices from multiple manufacturers via Windows.

In August 2025, ESET researchers published PromptLock, a cross-platform (Windows/macOS/Linux) proof-of-concept ransomware that uses a locally-run LLM to generate malicious scripts dynamically. It is relevant as evidence of the broader AI-malware trend, but it is a research PoC with no observed real-world attacks and is not ATM-specific.

The most disturbing fact:

76% of the malware detected in 2025 exhibits AI-driven polymorphism, and 37% shows evidence of optimization techniques using machine learning.

Every copy of the malware is different:
traditional detection signatures are no longer sufficient.

AI as an offensive multiplier

AI has not invented new attack categories. It has multiplied their scale, speed, and accessibility exponentially.

Four impact vectors are especially relevant to the ATM ecosystem:

Adaptive malware

Code that rewrites itself to evade defences. 76% of malware detected in 2025 is AI-polymorphic using ML optimisation techniques (Bitsight H1 2025).

Deepfake social engineering

44% of financial professionals report deepfakes in fraud schemes. Voice cloning with just 3 seconds of audio achieves 85% accuracy, directly relevant to staff managing remote ATM software updates.

Autonomous attack chains

Agentic AI systems execute full attack cycles, reconnaissance, payload, exfiltration autonomously. The window between a published vulnerability and active exploitation has compressed from weeks to hours.

Malware-as-a-Service (MaaS)

Tools such as WormGPT and FraudGPT allow actors with no technical skills to generate functional malware on demand, industrialising ATM crime.



Projected financial impact:

+1,210%

Growth in AI-enabled fraud in 2025, vs +195% for traditional fraud (Pindrop, 2026)

$40B

Projected AI fraud losses in the US by 2027 (Deloitte Center for Financial Services)



The weakest link

The untested software in production

Exploitation of vulnerabilities was the primary cause of 40% of attacks in 2025 (IBM X-Force 2026). Many of these vulnerabilities were not unknown zero-days: they were detectable flaws in software that made it into production without sufficient validation.

In the ATM ecosystem, three factors structurally exacerbate this risk:

Unpatched legacy OS

Many ATMs run Windows with monthly or quarterly update cycles and lack modern endpoint protections.

Vulnerable XFS middleware

The exact vector exploited by Ploutus. Its security validation is not always under the entity's direct control.

Short maintenance windows

The pressure for rapid deployments during off-peak hours prioritizes speed over validation.

The exposure window has closed

Offensive AI has drastically reduced the time between the disclosure of a vulnerability and its widespread exploitation. What used to take weeks now takes hours. AI systems automatically scan fleets of ATMs, identify vulnerable versions, and deploy payloads in fully automated cycles.

Software with known vulnerabilities that could once wait for a scheduled patch may now be compromised before the security team completes its impact assessment. Pre-production testing is the “only” way to close that window before it opens.

The real cost of not testing (properly)

Category	Impact
Direct fraud	Cash disbursed, write-offs
Operational downtime	+$300K/hour (ITIC)
Forensics and remediation	Hardware + equipment + time
Regulatory costs	Fines, audits, plans
Reputational damage	Loss of customers
Average cost per breach	**$5.56M (IBM 2025)**



> The question is no longer whether ATMs will be attacked.
>
> The question is whether the software running on them has been tested to withstand such an attack when the time comes.

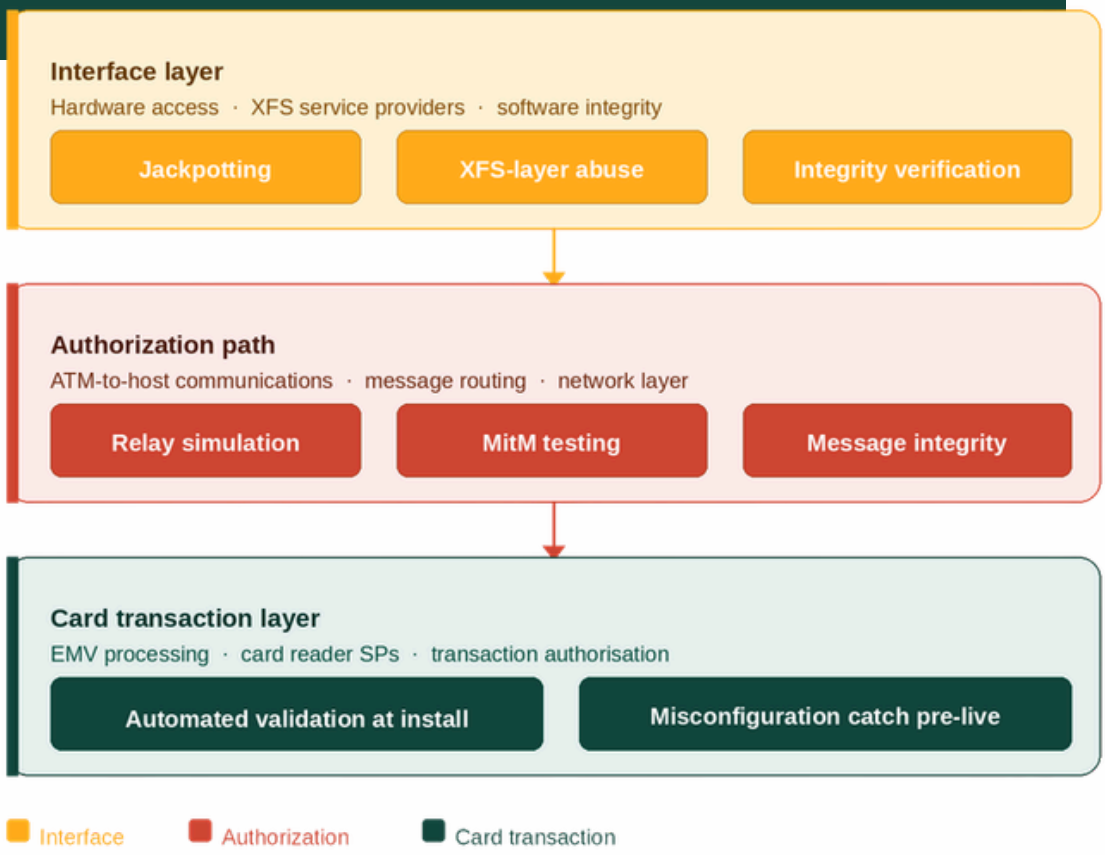
Mikel Aguirregabiria - Head of BussDev at Serquo

Where offensive testing fits

Serquo's simulation software lets a bank exercise its ATM software at scale, automatically, repeatedly, and before deployment.

Vigilis Defense adds an offensive layer on top of that simulation. Once the simulator confirms the software behaves as specified, Vigilis Defense attacks the same build the way a real adversary would, confirming that it contains no misconfiguration or weakness that could be turned into an unauthorized transaction. **Simulation proves the software is correct; offensive testing proves it is safe.**

Across the threat scenarios described above, that division maps directly onto the attack surface. Jackpotting and XFS-layer abuse — the vector exploited by Ploutus — are exercised at the interface level, alongside physical access controls and software-integrity verification.



The path to the authorization host is probed through relay-attack simulation, man-in-the-middle testing, and host-to-host message-integrity checks. And the card transaction layer is validated automatically at every new installation, catching the misconfigurations that quietly permit unauthorized transactions before a terminal ever goes live.

Just as important is what offensive testing can do that a specification-based simulator cannot: model attacks that do not yet have a signature. A simulator confirms that software meets known requirements; offensive testing asks the harder question: “can this build be broken by a technique that has not yet been catalogued?”.

In an environment where AI is compressing the gap between a vulnerability’s disclosure and its exploitation from weeks to hours, that question is no longer optional.

> “Automated verification tools are the surest way to guarantee the integrity of installed applications and that they are configured correctly.”
>
> *Zdravko Zdravkov, CTO & Lead Researcher, Vigilis Defense*



Pre-production testing as a strategic response

What effective testing should cover

SAST/DAST analysis	Static and dynamic vulnerability detection, with a special focus on the XFS layer and the modules that communicate with the banking host.
ATM penetration testing	Simulation of real attacks: XFS bypass, command injection, validation of physical-logical access controls.
Validation against the gold image	The FBI recommends continuously verifying that the deployed software does not deviate from the controlled reference image.
Integration testing	Validation of the interfaces between the ATM software and the authorization systems—the vector exploited by jackpotting.
Continuous DevSecOps	Automated security controls at every stage of the pipeline. Those who implement them save an average of $1.9 million per breach (IBM).

Serquo's products automate and scale this coverage while Vigilis Defense supplies the adversarial half and validate static and dynamic findings on a live rig to confirm real exploitability beyond scanner output.

Vigilis Defense runs the penetration tests — jackpotting and XFS-interface attacks, command injection, and physical-logical access review and, then, verifies software integrity against the gold image. It probes the authorization path with relay and man-in-the-middle attacks.

At deployment, Vigilis Defense's automated card-layer validation runs at every new installation, catching misconfigurations before they reach production.

Serquo's simulation confirms the software works; Vigilis Defense's offensive testing confirms it cannot be turned against the bank.

What effective ATM testing should cover

Testing area	Serquo Automation & scale	Vigilis Defense Adversarial half
SAST / DAST analysis	Automates static & dynamic scans Focus: XFS layer + banking-host modules	Validates findings on a live rig Confirms real exploitability beyond scanner
ATM penetration testing	Provides ATM rig & controlled environment Scales test execution across fleet	Jackpotting, XFS-interface attacks Command injection + access control review
Gold image verification	Automates integrity checks at every deploy FBI guideline: continuous reference audit	Adversarial validation of integrity controls Attempts bypass of image enforcement
Integration & authorization path	Instruments ATM-to-host interfaces Automated regression on auth flows	Relay and MitM attacks on auth channel Probes jackpotting vector end-to-end
Continuous DevSecOps	Automated controls across the pipeline Avg. $1.9M saved per breach (IBM 2025)	Threat modelling feeds pipeline controls Offensive findings close as pipeline gates

Serquo's simulation confirms the software works.

Offensive testing confirms it cannot be turned against the bank.

List of recommendations:

	Priority
1. Mandatory SAST/DAST testing prior to any ATM deployment	Critical
2. Establish and maintain a validated gold image for each ATM model	Critical
3. ATM penetration testing at least twice a year (focus on the XFS layer)	High
4. Integrate DevSecOps with automated controls in the pipeline	High
5. Continuous monitoring of software integrity in production	High
6. Staff training on deepfakes and AI-powered social engineering	Medium
7. Security audit of third-party software	Medium
8. Adopt Europol/NIST hardening guidelines specific to ATMs	Medium



About Serquo

Serquo is a technology company with over 30 years of experience. It specializes in financial device connectivity, payment solutions, web development, and offering efficiency solutions for ATM teams within financial institutions.

Serquo provides innovative tools that help banks and service providers streamline operations, automate testing, and accelerate development cycles for self-service technologies.

About Vigilis Defense

An offensive-security firm based in Sofia, founded in 2025, focused on the security of payment systems — POS and ATM terminals. We combine automated testing of card-transaction-layer vulnerabilities with hands-on offensive techniques, verifying that each build not only works as specified but also withstands the methods a real attacker would use.

Serquo and Vigilis Defense deliver end-to-end ATM security testing by combining automated simulation at scale with hands-on offensive techniques.

For more information, please, contact:
iinfo@serquo.com

Serquo.com
https://vigilisdefense.com