SILENT **HEISTS**

With infiltration efforts bolstered by AI, criminals increasingly favor under-the-radar data theft over outright disruption

There's a duality to cybercrime. Some attacks are loud, destructive, and impossible to miss. Yet some of the most dangerous breaches don't announce themselves. Attackers arrive quietly, move invisibly, and leave without a trace. By the time they're noticed, it's usually too late. This is the world of data exfiltration: Secretive digital heists in which thieves deliberately make no noise.

For modern cybercriminals, stealing data is profitable and elegant; it's more like a velvet glove than a wrecking ball. Exfiltration breaches each cost their victims an average of over USD 4.6 million, and the mean time to identify and contain a data breach is an astonishing 276 days.[‡] By the time a business realizes it has been compromised, it is often already too late to stem the bleeding.

The aftermath of public disruption news headlines, and crippled reputations can linger long after the attack. Who can forget the Ashley Madison breach, even a decade on, or the impact of Edward Snowden's downloaded cache of classified NSA materials?

Stolen credentials can expose critical systems, as seen in the 2024 Snowflake breach affecting over 160 companies

Exfiltration vs ransomware

Data exfiltration is not divorced from ransomware. Many ransomware attacks now include an extortion component, based on stolen data. In others, the high-profile part of the incident is something of a red herring to mask data theft. With ransomware fatigue growing—IBM reports that only 37% of 2025 victims opted to pay the ransom, down from 41% in 2024—criminal tactics are shifting toward the value of data.

Simultaneously, the regulatory costs to a business affected by a breach can be far more damaging than the temporary disruption caused by ransomware. Expenses for data mining reviews and notification obligations skyrocket in large breaches, and class action and settlement costs can also be crippling. In 2024 alone, approximately 2,000 data privacy lawsuits were brought to U.S. federal courts.[‡]

Two thirds of ethical hackers say they could start to exfiltrate data within five hours of beginning an attack—and **40%** suggest it would take less than two hours

Accelerated threats with AI

AI-driven malware and intrusion tools are making exfiltration faster, more targeted, and far harder to detect. Where traditional attacks might have relied on noisy trial-and-error, modern AI models can adapt in real time, learning the patterns of a victim's network to blend in with legitimate activity. AI-driven phishing campaigns with a machine learning component are becoming increasingly effective at gaining access, and AI-driven living-off-the-land (LOTL) techniques provide attackers with precision behavior within the perimeter.

Of **430 security incidents** handled by the UK's National Cyber Security Centre in 2024, **347 involved some level of data exfiltration**

The trajectory is clear: As ransomware's shock value declines, AI-enhanced exfiltration offers criminals a quieter, more profitable model. AI can even lead well-meaning employees to breach data unwittingly. Shadow AI tools, unapproved AI platforms embraced by workers seeking an edge in their tasks, are responsible for one in five breaches and add an average of USD 670,000 to breach costs for affected organizations.[‡]

The scale of the problem

Stolen credentials open the door to deeper systems, such as cloud platforms and third-party vendors. The 2024 Snowflake Breach compromised cloud data belonging to more than 160 companies, including AT&T and Santander, stealing customer records, government IDs, and sensitive corporate files[‡] from customer instances unprotected by multi-factor authentication.

Data theft is often about more than money or business disruption. At the nation-state level, exfiltration techniques are frequently used for critical infrastructure espionage. Addressing attendees at the 2024 Vanderbilt Summit, then FBI Director Christopher Wray warned of worrying findings from a honeypot experiment aimed at Chinese state actors: "It took the hackers all of 15 minutes to steal data related to control and monitoring systems, while ignoring financial and business-related information, which suggests their goals were even more sinister than stealing a leg up economically."

Staying one step ahead

Your strongest defense against AI-generated malware is Allowlisting. With ThreatLocker® Allowlisting, you control exactly what software can run in your environment—everything else, including unknown AI-generated threats, is blocked by default.

For an extra layer of protection, an agile, real-time endpoint detection and response (EDR) solution like ThreatLocker Detect can further harden defenses. Detect monitors behavior, analyzes activity in real time, and spots even the subtlest signs of adversarial activity. It acts immediately to stop malicious processes, block unauthorized applications, and prevent abnormal outbound connections before they become data leaks.

As AI-driven malware continues to evolve, organizations that combine Allowlisting as the frontline with EDR as an active defense layer will stay ahead of attackers and dramatically reduce the window for potential breaches. For an additional layer of confidence in EDR, organizations should also build the reassurance of remote monitoring and management (MDR) or security operations center (SOC) into their security portfolio, because there is no substitute for professional emergency assistance.

THREATLOCKER TIP

AI-powered defenses can still miss unknown AI-generated malware. Use Allowlisting to block AI-generated malware by default.