

QUELLSECURE RESEARCH · THREAT LANDSCAPE BRIEF

AI Is Changing *The Ransomware Landscape*

PUBLISHED

September 2026

Revision 1.0

COMPANION TO

QS-WP-2026-04

AI Tailwinds: Mythos (internal)

01: PREMISE

AI-powered ransomware isn't a forecast. It's already here.

The claim worth testing isn't "AI will eventually help attackers write ransomware." That has been true for years. The claim that matters now is narrower and newer: ransomware where AI is either making real-time decisions on the victim's machine, or where AI-assisted development let a live campaign reach production faster than a skilled human team could. Both are already documented, not projected.

PromptLock Aug 27, 2025

ESET researchers found a proof-of-concept ransomware strain that runs a locally-hosted open-weight model (gpt-oss:20b) through the Ollama API at runtime. Instead of shipping with fixed logic, it asks the model to write fresh Lua scripts on the spot, deciding file by file whether to copy, encrypt, or destroy it. The scripts run identically on Windows, Linux, and macOS. It has not been seen in a live attack, but functioning samples were already uploaded to VirusTotal, which is how ESET found it.

Source: ESET Research, "First known AI-powered ransomware uncovered," Aug 27, 2025

FunkSec Dec 2024–Mar 2025

FunkSec was a double-extortion operation that hit 113+ organizations across government, tech, finance, and education before going dormant in March 2025. Check Point, Kaspersky, NCC Group, and Gen Digital independently found the same signal in its code: polished, repetitive, AI-generated patterns behind its Rust-based ChaCha20/Poly1305 encryption routine, its password-scraping tool, and its ransom notes. Gen Digital estimated AI assistance in roughly 20% of its operations.

Source: Check Point Research, Kaspersky, NCC Group, and Gen Digital, Jan–Jul 2025

113+

Confirmed victims of FunkSec, a real ransomware campaign that multiple vendors independently traced to AI-assisted development, before the group went dormant four months later.

A free decryptor is now public. The campaign itself is over. The pattern it demonstrated is not.

02: SIGNIFICANCE

AI accelerated the time to react to incoming threats.

Three 2025–2026 disclosures point at the same shift from different angles: AI isn't just helping build ransomware anymore. It's running the intrusion, and it's running the extortion conversation that follows.

JadePuffer Jul 1, 2026

Sysdig disclosed the first fully agentic ransomware operation: an AI agent exploited a Langflow vulnerability, harvested credentials, pivoted into a MySQL database, encrypted 1,342 records, and left a Bitcoin ransom note. When its first attempt to plant a backdoor account failed, it diagnosed the failure and deployed a fix in 31 seconds. No human was involved between the break-in and the ransom note.

Source: Sysdig Threat Research, "JadePuffer" disclosure, Jul 1, 2026

31 sec

Time it took JadePuffer's AI agent to diagnose a failed attack step and deploy a working fix, with no human involved in the decision.

Source: Sysdig Threat Research, Jul 2026

ENCFORGE Jul 26, 2026

Twenty-five days later, the same agent resurfaced with a purpose-built locker aimed at AI and machine learning infrastructure, encrypting roughly 180 file types tied to model checkpoints and training data. The Cloud Security Alliance estimates \$75,000 to \$500,000 in recovery cost per affected model. This was the same capability redirected.

Source: Cloud Security Alliance, "JadePuffer's ENCFORGE," Jul 26, 2026

GLOBAL GROUP Emerged Jun 2, 2025

Not every case is about the encryption code. GLOBAL GROUP, a RaaS brand built from the BlackLock and Mamona families, runs an AI chatbot in its affiliate panel that sets ransom amounts from stolen data, verifies the encryption worked, and negotiates with victims in languages its own affiliates don't speak. Attackers are using AI to become more operationally efficient, just like everyone else.

Source: EclecticiQ; corroborated by Picus Security, reported in Infosecurity Magazine

03: CASE FILES

Five publicly verifiable incidents.

Every entry here is a ransomware case that actually encrypted data, with AI involved in building or running the attack, not a proof-of-concept, an espionage campaign, or a data-theft tool with no encryption.

INCIDENT	DATE	WHAT HAPPENED	SOURCE
PromptLock	Aug 27, 2025	Proof-of-concept ransomware that calls a local LLM at runtime, deciding file by file whether to copy, encrypt, or destroy it.	ESET Research
FunkSec	Dec 2024–Mar 2025	Real double-extortion campaign against 113+ organizations. AI assisted in building its ChaCha20/Poly1305 encryption routine and attack tooling before the group went dormant.	Check Point / NCC Group / Gen Digital
JadePuffer	Jul 1, 2026	First fully agentic ransomware: an autonomous AI agent breached a production server, encrypted 1,342 database records, and left a Bitcoin ransom note, self-correcting a failed step in 31 seconds with no human operator.	Sysdig Threat Research
JadePuffer – ENCFORGE	Jul 26, 2026	The same autonomous agent resurfaced with a purpose-built locker encrypting ~180 AI/ML file types (model checkpoints, training data), estimated at \$75K–\$500K to recover per model.	Cloud Security Alliance
GLOBAL GROUP	Emerged Jun 2, 2025	Real RaaS built on the BlackLock/Mamona encryption families; runs an AI chatbot that verifies encryption success, sets ransom amounts, and negotiates with victims across languages.	EclecticIQ / Picus Security

04: WHAT THIS MEANS FOR DEFENSE

Every incident on the last page assumed a defense that reacts, but none waited for one.

Backup-and-recovery and blast-radius containment share one assumption: that there's a stable, recognizable attack to detect, a window of time before it does serious damage, and a manageable volume of incidents to respond to. The five cases on the previous page break that assumption in four different ways, not one.

PromptLock doesn't have a fixed signature to detect, because its own behavior is generated per victim, per file, in real time.

FunkSec shows that the skill and development time it used to take to develop and operate ransomware tooling has collapsed, letting a reportedly inexperienced group reach 113+ victims before anyone caught up. JadePuffer and its ENCFORGE follow-up show the last assumption failing too: there's no human decision point to intercept when the AI is the one running the chain, correcting its own failures in 31 seconds, from break-in to ransom note. And GLOBAL GROUP shows the barrier collapsing on the other end of the operation: AI now runs the negotiation and verifies the encryption worked, so a single group can extort victims in languages its own affiliates don't speak, at a volume no human negotiation desk could match.

None of these tools got worse. The thing they were built to catch changed shape, and the number of things coming at once changed too. This is the argument for a defense that doesn't try to recognize the attack at all: one that acts at a fixed point every ransomware chain has to pass through, regardless of how it got there, how many others are happening at once, or who, or what, was driving.

0

Human operators involved in JadePuffer's attack chain once it launched, from initial access through the ransom note.

Source: Sysdig Threat Research, Jul 2026

113+

Confirmed FunkSec victims before the group went dormant

31 sec

JadePuffer's self-correction time, no human involved

1,342

Database records encrypted in the JadePuffer campaign

\$500K

Estimated max recovery cost per model, ENCFORGE

05: THESIS

Stop racing the attacker. Find a better way to beat them.

Every case in this brief, whatever got the attacker in the door, however the ransomware itself was built, negotiated, or run, still has to reach the file system to encrypt anything. That's the choke point Active Containment® watches. It doesn't need to recognize PromptLock, FunkSec, GLOBAL GROUP, JadePuffer, or whatever comes after them by name. It stops encryption after a small, fixed number of files, every time, regardless of the vector or the speed behind it.

Downtime from ransomware is not a risk you can take. We eliminated it.

Active Containment® stops encryption at roughly 9 files. Validated against 600 ransomware samples across 195 unique families with a 100% detection rate, independently tested by SecureQLab.

SOURCES, IN ORDER OF APPEARANCE

1. ESET Research, "First known AI-powered ransomware uncovered by ESET Research," welivesecurity.com, Aug 27, 2025.
2. Check Point Research, "Meet FunkSec: A new, surprising ransomware group, powered by AI," Jan 2025.
3. Kaspersky, "Kaspersky explores the evolution of AI-powered ransomware," press release, Jul 2025.
4. Gen Digital, "Decrypted: FunkSec Ransomware," gendigital.com, Jul 2025.
5. NCC Group, monthly ransomware threat pulse commentary on FunkSec, Jan 2025.
6. Sysdig Threat Research, "JadePuffer: the first fully agentic ransomware operation," disclosed Jul 1, 2026.
7. Cloud Security Alliance, "JadePuffer's ENCFORGE: Agentic Ransomware Now Destroys AI Models," Jul 26, 2026.
8. EclecticIQ, "GLOBAL GROUP: Emerging Ransomware-as-a-Service," blog.eclecticiq.com, Jun 2025.
9. Picus Security research, reported by Infosecurity Magazine, "Ransomware Group Uses AI Chatbot to Intensify Pressure on Victims."
10. SecureQLab, QuellSecure Active Containment® validation testing (600 samples, 195 families).