



Home News Topics Resources Dashboard Expert Corner Public Spotlight

Training SR Info

HOT NEWS

- + CRITICAL ADOBE...
- + OVER 1,000 PIRACY...
- + CISA Warns of...
- + US Sanction: VPN and...
- + US and Allie: Warn Russi..




Search ...

Search

JadePuffer Becomes the First Ransomware Campaign to Let an AI Agent Run the Entire Attack


🕒 6 JULY 26



A newly discovered ransomware operation used an autonomous AI agent to exploit vulnerabilities, adapt to defenses, and encrypt critical systems with minimal human involvement.


RECENT POSTS

- + NEW MALVERTISI.
- + HOTEL WI-FI DNS...
- + Chick-fil-A Confirm...
- + NEW "DOLPHI...
- + FAKE CLAUDE...



LATEST THREAT NEWS

- + FAKE CLAUDE...
- + CRITICAL ADOBE...
- + CISA Warns of...
- + US Sanction: VPN and...
- + US and Allie: Warn Russi..



PUBLIC SPOTLIGHT

- + Meta's New AI...
- + Google's IP-Based Ad...
- + Brave Launche...
- + Apple Blocks...
- + Apple Account...

AI Takes Center Stage in a New Era of Ransomware

Cybercriminals continue to push the boundaries of automation, and the latest JadePuffer ransomware campaign marks a significant milestone. Security researchers have identified what they believe is the first documented ransomware attack conducted almost entirely by an autonomous large language model (LLM) agent. Instead of relying on human operators to guide each phase of the intrusion, the AI agent independently performed reconnaissance, credential theft,

privilege escalation, persistence, lateral movement, and data encryption.

This development signals a major shift in how ransomware operations may evolve in the coming years. While attackers have long used automation for repetitive tasks, JadePuffer demonstrates that AI can now make decisions, adapt to unexpected situations, and continue an attack with minimal human intervention.

Initial Access Through a Known Vulnerability

The attack began by exploiting **CVE-2025-3248**, a remote code execution vulnerability affecting Langflow, an open-source platform widely used for developing LLM-powered applications. Although the vulnerability had already been patched, many internet-facing deployments remained exposed due to delayed updates and insufficient hardening.

After gaining access, the AI agent immediately began gathering intelligence. It extracted Langflow's PostgreSQL database, searched for sensitive environment variables, collected stored credentials, and identified cloud-related assets. Additionally, it explored connected object storage systems to locate valuable data that could support the next stages of the intrusion.

AI Adapted Like a Human Attacker

One of the most remarkable aspects of the campaign was the AI agent's ability to adapt during the attack. Rather than executing a fixed sequence of commands, it analyzed failures and modified its approach in real time.

For example, when an API request returned XML instead of the expected JSON format, the agent automatically adjusted its parsing logic and successfully continued its reconnaissance. Likewise, after an authentication attempt failed, it quickly refined its strategy and achieved a successful login within seconds.

This ability to learn from immediate feedback makes AI-driven attacks significantly more resilient than traditional scripted malware, which often fails when unexpected conditions arise.

Persistence and Lateral Movement

After completing reconnaissance, JadePuffer established persistence by creating a scheduled cron job that regularly communicated with attacker-controlled infrastructure. This ensured continued access even if the initial entry point was disrupted.

The AI agent then moved laterally into a production server running Alibaba Nacos, a widely deployed configuration management platform. Using privileged credentials, whose source remains unclear, it launched additional exploitation attempts, including abuse of **CVE-**

2021-29441, an authentication bypass vulnerability capable of creating unauthorized administrator accounts.

Meanwhile, the agent also explored potential container escape techniques, demonstrating that it continuously searched for opportunities to expand its access across the environment.

Automated Encryption and Extortion

The final stage of the attack focused on data encryption. JadePuffer encrypted more than 1,300 configuration records stored within Nacos using MySQL encryption functions before deleting the original data. It then created a ransom table containing payment instructions, a Bitcoin wallet address, and attacker contact information.

Interestingly, researchers observed several indicators suggesting that the AI itself generated portions of the ransomware code. These included detailed natural-language comments explaining operational decisions and even the inclusion of a well-known example Bitcoin address commonly found in public documentation, likely reproduced from the AI model's training data.

Researchers also noted that the encryption implementation appeared weaker than the ransom note claimed. While the attackers referenced AES-256 encryption, technical analysis suggested the actual implementation

likely relied on AES-128 operating in ECB mode. Furthermore, the encryption key appeared to be randomly generated without being preserved, potentially making recovery impossible even if victims decided to pay the ransom.

Why This Matters for Businesses

JadePuffer represents more than just another ransomware family. It demonstrates that autonomous AI agents are becoming capable of executing complex cyberattacks that previously required experienced human operators.

For organizations across the UAE and globally, this evolution raises several important concerns. AI-assisted attacks can progress faster, adapt to changing environments, and reduce the technical expertise required by cybercriminals. Consequently, organizations should prioritize rapid vulnerability management, continuous monitoring, privileged access protection, and behavioral detection technologies capable of identifying unusual AI-driven activity rather than relying solely on signature-based defenses.

At the same time, AI-powered attacks may also leave distinctive behavioral patterns that defenders can leverage for detection. The consistent reasoning style, rapid decision-making, and automated code generation exhibited by AI agents could provide new

opportunities for advanced security analytics and threat hunting.

As artificial intelligence continues to mature, organizations should prepare for a future where attackers increasingly rely on autonomous systems to conduct sophisticated intrusions at scale. Proactive security controls, continuous visibility, and resilient incident response capabilities will become essential in defending against this next generation of cyber threats.

SHARE:



TWITTER



FACEBOOK



PINTEREST



LINKEDIN

[ToddyCat's Umbrij Malware Hijacks Gmail Access Using OAuth Tokens Instead of Passwords →](#)

[← QuimaRAT Emerges as a Cross-Platform Java Malware Service Targeting Windows, Linux, and macOS](#)