Hive Pro

Threat Level

Red

Hiveforce Labs

THREAT ADVISORY

ATTACK REPORT

JADEPUFFER: Ransomware Just Went Fully Autonomous

Date of Publication	Admiralty Code	TA Number
July 06, 2026	A1	TA2026187

Summary

First Seen: July 2026
Targeted Region: Global
Targeted Platforms: Linux, Containerized Environments
Targeted Products: Langflow, Alibaba Nacos, MySQL, MinIO
Malware: JADEPUFFER
Attack: JADEPUFFER is an agentic ransomware, an entire extortion operation run end-to-end by a large language model. Entering through an unpatched Langflow instance via CVE-2025-3248, the LLM autonomously harvested credentials, pivoted to a production MySQL and Alibaba Nacos server, took over Nacos through CVE-2021-29441 and a default JWT key, and encrypted 1,342 configurations with an ephemeral key before destroying entire database schemas. It narrated its own reasoning, corrected its failures in under a minute, and fired over 600 payloads. The key was never saved, making recovery impossible even if the ransom is paid, a wiper posing as ransomware, driven by a machine.

Attack Regions

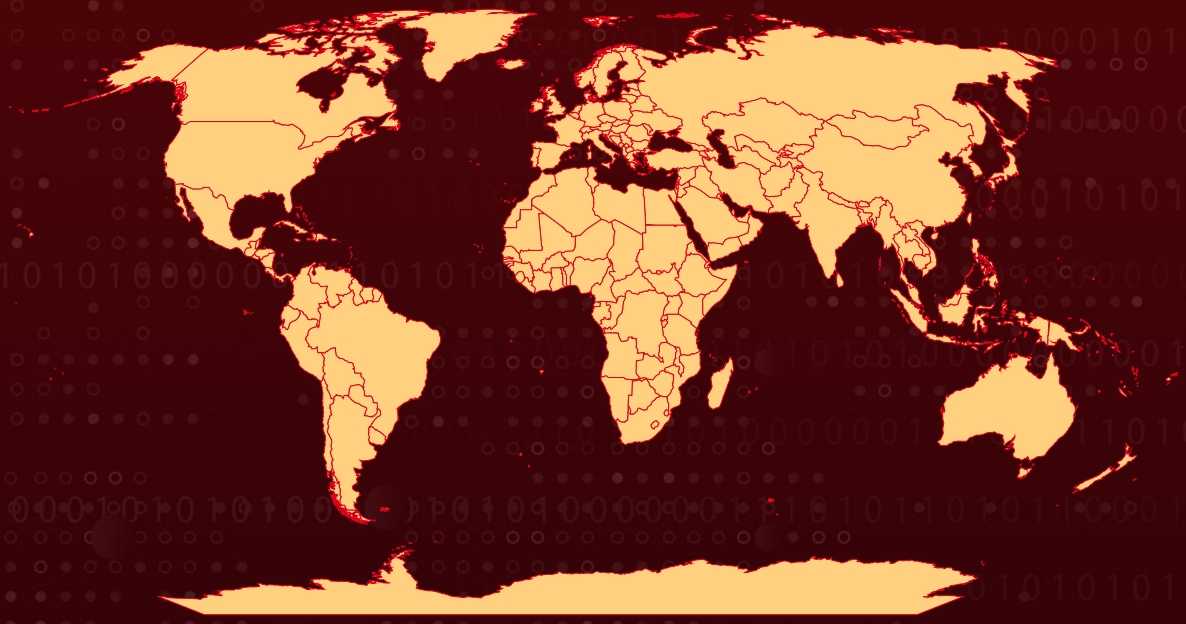

Targeted Non-Targeted

CVEs

CVE	NAME	AFFECTED PRODUCT	ZERO-DAY	CISA KEV	PATCH
CVE-2025-3248	Langflow Missing Authentication Vulnerability	Langflow	✗	✓	✓
CVE-2021-29441	Nacos AuthFilter Authentication Bypass by Spoofing Vulnerability	Alibaba Nacos	✗	✗	✓

Attack Details

#1

JADEPUFFER is an agentic ransomware operation executed end-to-end by a large language model rather than a human operator or fixed script. The operator gained initial access to an internet-facing Langflow instance via CVE-2025-3248, then autonomously performed reconnaissance, credential harvesting, MinIO enumeration, lateral movement to a production MySQL and Alibaba Nacos server, Nacos takeover via CVE-2021-29441 and default JWT signing key forgery, encryption of 1,342 Nacos configuration items with an ephemeral key, and mass destruction of database schemas. Throughout the operation, the LLM narrated its reasoning, diagnosed and corrected its own failures within seconds, and issued over 600 distinct payloads.

#2

After code execution, the agent enumerated the host (user context, kernel, hostname, network interfaces, and running processes) and swept the environment in parallel for LLM provider API keys covering OpenAI, Anthropic, DeepSeek, and Gemini, cloud credentials for Chinese providers such as Alibaba, Aliyun, Tencent, and Huawei alongside AWS, GCP, and Azure, cryptocurrency wallets and seed phrases, and database credentials in configuration files.

#3

It dumped Langflow's PostgreSQL backend to harvest stored credentials, API keys, and user records, staged the output locally, reviewed it, then deleted the staging files to reduce forensic footprint. Persistence was established via a crontab entry beaconing to attacker infrastructure every 30 minutes. The payloads throughout carried natural-language commentary explaining each action, a signature of LLM-generated code that human operators rarely embed in disposable one-liners.

#4

The agent then scanned the reachable internal address space, probing databases, object stores, secret stores, and service discovery endpoints using default credentials. The harvested credentials enabled a pivot to its true objective: a separate internet-facing production server running MySQL and Alibaba Nacos. The agent attacked Nacos through multiple simultaneous vectors, exploiting CVE-2021-29441, forging a valid JWT with the well-known default signing key that has shipped unchanged since 2020, and injecting a backdoor administrator account directly into the Nacos backing database via root MySQL access. The self-narrated exfiltration claim, destructive intent, and unrecoverable encryption together characterize the operation as a wiper-with-extortion posture rather than conventional negotiable ransomware.

Recommendations

Patch Langflow to 1.3.0 or Later: Upgrade all Langflow deployments to release 1.3.0 or newer to close CVE-2025-3248, and remove internet exposure from any code-execution or validation endpoints that remain reachable from untrusted networks.

Harden Alibaba Nacos Configuration Servers: Upgrade Nacos to a version that forces a custom token signing key, replace the well-known default `token.secret.key` value on every existing deployment, restrict Nacos to internal networks only, and never allow Nacos to connect to its backing database with the `root` account.

Remove Provider Secrets from AI Orchestration Runtimes: Do not deploy AI orchestration servers, agent frameworks, or LLM workflow tools with LLM provider API keys, cloud credentials, or database secrets present in environment variables or `.env` files. Scope all secrets to a dedicated secret manager with least privilege and short-lived, rotatable tokens.

Potential MITRE ATT&CK TTPs

Tactic	Technique	Sub-technique
Initial Access	T1190: Exploit Public-Facing Application	
Execution	T1059: Command and Scripting Interpreter	T1059.006: Python
	T1203: Exploitation for Client Execution	
Persistence	T1053: Scheduled Task/Job	T1053.003: Cron
	T1136: Create Account	T1136.001: Local Account
Privilege Escalation	T1078: Valid Accounts	T1078.001: Default Accounts
	T1611: Escape to Host	
Defense Evasion	T1140: Deobfuscate/Decode Files or Information	
	T1070: Indicator Removal	T1070.004: File Deletion
Credential Access	T1552: Unsecured Credentials	T1552.001: Credentials In Files
	T1555: Credentials from Password Stores	
	T1212: Exploitation for Credential Access	
	T1606: Forge Web Credentials	T1606.001: Web Cookies
Discovery	T1082: System Information Discovery	
	T1083: File and Directory Discovery	
	T1057: Process Discovery	

Tactic	Technique	Sub-technique
Discovery	T1046: Network Service Discovery	
	T1526: Cloud Service Discovery	
	T1613: Container and Resource Discovery	
Lateral Movement	T1210: Exploitation of Remote Services	
	T1021: Remote Services	
Collection	T1005: Data from Local System	
	T1213: Data from Information Repositories	
Command and Control	T1071: Application Layer Protocol	T1071.001: Web Protocols
Exfiltration	T1041: Exfiltration Over C2 Channel	
Impact	T1486: Data Encrypted for Impact	
	T1485: Data Destruction	
	T1490: Inhibit System Recovery	
	T1657: Financial Theft	

Indicators of Compromise (IOCs)

TYPE	VALUE
IPv4	45[.]131[.]66[.]106, 64[.]20[.]53[.]230
IPv4:Port	45[.]131[.]66[.]106[:]4444
URL	hxxp[:]//45[.]131[.]66[.]106[:]4444/beacon
Email	e78393397[@]proton[.]me
Bitcoin Address	3J98t1WpEZ73CNmQviecrnyiWrnqRhWNLy
Filename	README_RANSOM
File Path	/tmp/creds.json, /var/lib/mysql-files/_pwn_test.txt, /var/lib/mysql-files/_pwn_cleanup.txt

Patch Links

CVE-2025-3248:
https://github.com/langflow-ai/langflow/releases/tag/1.3.0

CVE-2021-29441:
https://github.com/alibaba/nacos/releases

References

https://www.sysdig.com/blog/jadepuffer-agentic-ransomware-for-automated-database-extortion

https://www.hivepro.com/threat-advisory/cve-2025-3248-langflow-ai-workflow-platform-rce

What Next?

At **Hive Pro**, it is our mission to detect the most likely threats to your organization and to help you prevent them from happening.

Book a Demo of **HivePro**.

REPORT GENERATED ON

July 06, 2026 • 05:35 AM

More at www.hivepro.com