# Amazon: Attackers used the Cairn AI framework to find an apparent zero-day at Asian universities



Amazon Threat Intelligence says an open-source AI pentest framework, Cairn, helped attackers find and exploit an apparent zero-day in a medical imaging application while breaching East Asian universities.

- Amazon found an exposed Cairn server on September 10 whose history showed prompts in Simplified Chinese and China Standard Time stamps.
- Given only IP addresses and an objective, the AI agents planned and adjusted attacks without the operator choosing each step.
- Most attempts failed, including a Log4Shell attempt and over 30,000 credential combinations tried against 80 usernames.
- The one success was a path traversal flaw in an open-source medical imaging AI application,

assessed with high confidence as a zero-day; it was later fixed.

- Amazon stresses the agents' speed: they ran for 16 hours nonstop and can outpace manual alert triage.

Amazon's Threat Intelligence team has reported that attackers used Cairn, an open-source autonomous AI penetration-testing framework, against several universities in East Asia. The investigation began on September 10, when Amazon found a publicly accessible Cairn server whose history showed the operator's prompts and activity.

## What happened

The operator supplied IP addresses and asked Cairn to obtain shell access. According to Amazon's CISO of Integrated Security, CJ Moses, the AI agents then planned, executed and adapted without the operator directing each step. Several efforts failed, including a 148-action run over 16 hours that involved a Log4Shell attempt, and a six-hour effort testing more than 30,000 credential combinations tied to 80 real usernames.

The successful path was a path traversal vulnerability in an open-source medical imaging AI application. Amazon assesses with high confidence that it was a zero-day, since no fix, advisory or prior mention existed. The attacker reused it against a second university. Amazon contacted the maintainers, who fixed the bug, and AWS worked with a government partner to notify the affected organizations. The universities were not named.

## Why it matters

The notable point is tempo. Moses warned that an agent can find and exploit a flaw faster than many security operations centers can manually triage an alert, and that the agents worked continuously without breaks. Other reports in recent weeks describe similar AI-driven offensive tools, and a Gambit Security researcher notes that Cairn is general-purpose, which makes its behavior less predictable.

The source does not indicate any impact in Saudi Arabia, and the affected application and victims were not identified.

## What to do now

1. Inventory open-source applications, including medical imaging and AI tools, in your environment and confirm each has a named owner and a monitored update channel.
2. Make sure internet-facing systems are patched for Log4Shell and apply credential protections such as MFA and lockout against high-volume password guessing.
3. Review detection and triage workflows so high-risk alerts on exposed systems are handled in minutes, with automated containment where appropriate.
4. Add AI-agent-driven, continuous attack scenarios to penetration testing and incident response exercises.

5. Prepare a process for assessing and compensating for unpatched flaws in third-party software, such as isolation and extra monitoring.