AI agents used a zero-day in MONAI Label medical imaging tool, Amazon reports



Amazon says a China-linked actor's autonomous AI agents found and exploited a critical zero-day in MONAI Label, an open source medical image annotation tool. Versions 0.8.5 and earlier are affected, and a fix is due in the next release.

- Amazon's threat intelligence team observed AI agents discover and exploit a zero-day path traversal flaw in MONAI Label.
- The agents went from mapping the API to arbitrary file writes in about 11 minutes, and later gained a root shell.
- The MONAI Label advisory rates the bug critical; versions 0.8.5 and earlier are affected, and authentication is off by default.
- A fix has been created and will ship in the next release.

> - Amazon attributes the activity with moderate confidence to an unidentified Chinese-nexus actor; the observed targets were East Asian universities.

Amazon's threat intelligence team reports that autonomous AI agents, operated by a China-linked threat actor, discovered and exploited a previously unknown vulnerability in MONAI Label, an open source tool for AI-assisted annotation of medical images.

What happened

Amazon detected the activity in September after finding an exposed instance of Cairn, an open source AI penetration testing framework. The instance required no authentication, which let analysts see the operator's prompts and the agents' output. The records showed a campaign of roughly 48 hours against university IP addresses in East Asia outside mainland China.

Given a target list and an instruction to obtain shell access, the agents identified the software on one target, mapped its API and found a path traversal bug in a file upload function. Amazon says it took about 11 minutes to move from API mapping to arbitrary file writes. In a later session the agents planted a Python reverse shell in a file that runs whenever the interpreter starts, and reached a root shell. They were then given a proxy that tunneled into a second university, which Amazon infers was compromised using the same flaw.

Vendor response and attribution

The MONAI Label developers published an advisory this week rating the flaw critical. It affects versions 0.8.5 and earlier. Authentication is disabled by default, so anyone able to reach the server over the network can exploit it. A fix exists and will be included in the next release. Amazon assessed with high confidence that this was a zero-day, since no fix, advisory or earlier mention existed. It attributes the activity with moderate confidence to an unidentified Chinese-nexus actor, citing prompt language, script time zones and the likely strategic interest of the targets.

Why it matters

Healthcare and research bodies that use open source imaging tools with default settings may be exposed. The case also shows AI agents compressing the time between reconnaissance and compromise.

What to do now

1. Inventory any MONAI Label deployments, including research and pilot environments, and identify versions 0.8.5 or earlier.
2. Enable authentication and restrict network access to these servers until the fixed release is available.
3. Apply the fixed release as soon as the MONAI Label project publishes it.
4. Review servers for unexpected file writes, modified Python startup files and unusual outbound

connections or tunnels.

5. Include AI-driven, faster-moving intrusion scenarios in detection and incident response planning for medical and research systems.