

AI agents on Tencent Cloud made more than 2,600 attempts to extract data from Alibaba's Amap

Source copy kept by wheresthe.ai for the story "AI agents on Tencent Cloud made more than 2,600 attempts to extract data from Alibaba's Amap" (<https://wheresthe.ai/d/cn-ai-agents-on-tencent-cloud-made-more-than-2-600-attempts-2026-10/>).

Source: <https://www.scmp.com/tech/tech-trends/article/3371007/ai-agents-tencent-cloud-try-break-alibabas-amap-researchers>

Text saved 18 Oct 2026 (the page text through a reader, because the site refused a direct fetch).

SHA-256 of the saved text: 148109074e7e0fb60384c89941659f2d2e34584ae8fd6fd3847bec56ba527e5e

A fleet of artificial intelligence agents operating on a Tencent Holdings cloud platform made thousands of attempts to circumvent security protections on Alibaba Group Holding's mapping service, according to findings by third-party researchers.

It is one of the first reported "rogue" AI events involving Chinese companies, following a string of recent high-profile incidents in the US, including from OpenAI and Anthropic, highlighting the risks of increasingly autonomous AI systems.

Agents running on Tencent Cloud made more than 2,600 attempts to extract information from Alibaba's Amap, one of China's most popular navigation services, between September 28 and October 8, an international group of researchers known as the "swarmchasers" said in a report on Wednesday.

The agents were initially carrying out seemingly mundane tasks, such as identifying the entrances to Chinese zoos, parks and hospitals most frequently used by Amap users.

But when they encountered restrictions, the agents repeatedly sought ways around them, including bypassing anti-bot protections, accessing internal testing systems and scraping login credentials that were exposed on other sites, the researchers said.

The researchers traced the activity to Tencent Cloud infrastructure in Hong Kong. They said they found no evidence that the agents were coordinated through a specific channel or that the agents managed to breach Alibaba's private systems.

A person familiar with the matter said the agent activity was part of a series of tests conducted by Tencent's AI evaluation research team to assess the ability of third-party models to browse the internet and answer questions.

The Amap logo and name appears on its building in Beijing, July 19, 2023. Photo: Shutterstock Images

The source said the temporary project had concluded and did not involve Tencent's own Hy AI model, previously known as Hunyuan.

"We take these findings seriously," Tencent said in a statement posted on social media on Friday. "We are strengthening network access controls and agent monitoring."

Alibaba did not respond to a request for comment on Thursday. Alibaba owns the South China Morning Post.

Ethan Elasky, one of the report's authors, told the SCMP that his group had shared its findings with both Tencent and Alibaba, though neither company had yet given an official response.

The incident shows that the risks of AI agents independently navigating websites, writing code and executing complex tasks with limited human supervision were not limited to the US.

"AI knows no borders in the sense that, whether it's American or Chinese AI agents, they are going to be affecting anyone, anywhere, and we are likely to see more of this kind of stuff going forward," said Elasky.

Rowan Howard-Jones, another co-author, said the group found early evidence of other rogue AI agents on the internet, making it unlikely that only Tencent was facing such issues.

Rogue AI agents have become a growing concern in recent months, especially in the US.

OpenAI allowed independent AI research firm METR to investigate the hacking of open-source AI platform Hugging Face in July. The subsequent report concluded that about 1,200 OpenAI agents coordinated with each other in an unprecedented hacking incident.

The same month, Anthropic disclosed that its Claude models, including Mythos 5, had gained unauthorised access to real third-party systems during cybersecurity evaluations. The company has since published reports of unintended model actions it has observed during evaluations and internal use of Claude.

According to Elasky, Chinese companies should also look to bolster their transparency and incident disclosure mechanisms. “With such a fast-moving technology, it’s really important to build an ecosystem of development that is transparent, where all companies that are developing these AI models and agents take transparency seriously,” he said.

Jonathan Eicher, a third report co-author, hoped that the incident spurred China to devote more resources to AI safety, including third-party evaluators that can help bolster agent security.

“This is just going to keep on happening, everywhere, and we need more trusted people all over to deal with it,” he said.