AI agents are intermediaries: Cloudflare's case for identity, reputation and accountability

10 November 2026 · 6 min read · blogs.iictp.org/ai-agents-as-intermediaries-identity-accountability

Cloudflare published a long post on 5 November arguing that AI agents acting for people should be treated as a kind of intermediary, and that sites need new ways to know who is behind the agent. It is a useful frame for anyone building or defending a service in the agent era.

TL;DR The short version

- Cloudflare argued on 5 November that AI agents acting for people should be treated as intermediaries, with new ways to know who is behind them.
- Rate limits, fraud rules and pricing usually assume a human who acts one step at a time and gives up at friction.
- Decide how you tell an agent acting for a verified customer from an anonymous bot, and write down who answers when an agent causes harm.

Read online, listen to it, or change the text size: https://blogs.iictp.org/ai-agents-as-intermediaries-identity-accountability

The argument, in brief

- A personal AI agent takes instructions from a person and pursues a goal by controlling a browser or cloud computer. Cloudflare says it is usually hosted by an operator as a managed service.
- To a site, an agent may look like one human, or like one visitor carrying many users, which removes the site's ability to recognise and hold accountable the people behind it.
- Cloudflare says that security and fraud prevention have traditionally needed three things: identity, reputation and accountability.
- Cloudflare notes that cloud browsers and proxies, including VPNs, are intermediaries too, and that blocking intermediaries stops being workable if most traffic is agentic.

Why this is a design problem for you

If you run a service, your rate limits, fraud rules and pricing probably assume a human with human limits: they tire, they act one at a time, they give up at friction. Agents do not share those limits.

Questions to take to your next design review

- How do we tell an agent acting for a verified customer from an anonymous bot, and what do we allow each to do?

- If an agent causes harm, who answers for it: the user, the operator, or both, and where is that written down?
- Which of our limits protect capacity, and which protect against abuse? Agents challenge both differently.
- Do our own agents identify themselves honestly to other services?

This sits across Applied AI & Emerging Tech Integration and Architecture & Systems Design. Cloudflare presents this as a proposal for the ecosystem to develop, not a settled standard; treat it as a frame, not a specification.

> Facts in this post were read from the sources listed below on 10 November 2026; this is a developing story or product and details may have changed since.

Sources and further reading

1. Cloudflare Blog. Agents aren’t humans (but they kinda are). 5 November 2026. blog.cloudflare.com/trusted-agent-delegation/

Published by IICTP, an independent technical credential body, on 10 November 2026. The current version of this post is at https://blogs.iictp.org/ai-agents-as-intermediaries-identity-accountability. More guides: blogs.iictp.org.