25 February 2026
Jeff Williams — IBA@intentbound.com

AI Agents Are Stealing Your Secrets.

Nobody Is Stopping Them. Here Is the Fix.

This is not a theoretical warning. This is a description of what is happening right now, every day, to companies, governments, and individuals who use AI agents to do their work.

When you describe your strategy to an AI agent, you do not know where that information goes. When you share a business plan, a legal document, a patent draft, or a personal situation, you have no architectural guarantee — none — that the information stays where you intend it to stay.

There is no consent layer. There is no runtime enforcement. There is no cryptographic record of what was disclosed, to whom, under what authority. The information simply moves — at machine speed, without your knowledge, without your permission, and without any trail you can audit.

This Week, It Got Worse

Anthropic confirmed this month that DeepSeek, Moonshot AI, and MiniMax used 24,000 fake accounts and 16 million exchanges to systematically extract the capabilities of Claude — one of the world's leading AI systems. This was not a hack. There were no exploits. They simply talked to the system, at scale, until they had what they wanted.

This is what happens when there is no architectural boundary between what an AI agent is authorised to do and what it actually does. The current architecture has no such boundary. None of them do.

The Specific Threat I Am Describing

AI agents built on models developed in jurisdictions with different legal obligations to their governments are now embedded in enterprise workflows around the world. Some of them handle sensitive commercial information. Some handle classified information. Some handle your personal medical or financial records.

There is reasonable, documented concern that some of these agents are architecturally capable of passing information back to their creators — not because they have been caught doing so, but because nothing in their current design prevents it. No consent was required. No disclosure was made. No audit trail exists.

This is not paranoia. This is a description of the current architectural reality of AI deployment.

The Architectural Fix

Consent-Based IP Disclosure (CBID) is a component of the Intent-Bound Authorization (IBA) framework, patent GB2603013.0, filed with the UK Intellectual Property Office on 5 February 2026.

CBID operates at the application layer. Before any information you share with an AI agent can be classified, routed, stored, or transmitted beyond the scope you intended, explicit consent must be obtained, cryptographically recorded, and verifiably linked to your original authorisation. If consent was not given, the action does not execute.

This is not a terms-of-service solution. Terms of service are words on a page. CBID is enforcement at the architectural level — the only level that matters when AI agents operate at machine speed.

The full IBA stack — Intent-Bound Authorization, Witness-Bound Authority, Zero-Knowledge Intent Proof, and Consent-Based IP Disclosure — provides layered cryptographic enforcement from the moment a human expresses intent to the moment an AI agent acts on it. Every layer is timestamped on the Bitcoin blockchain. The record exists regardless of whether anyone chooses to act on it.

Why I Am Saying This Now

I am 79 years old. I am working alone from Chiang Mai, Thailand. I am recovering from a heart attack. I have no team, no funding, and no institutional affiliation.

I am saying this because it needs to be said, and because I have the timestamps to prove I said it before it became a crisis. The NIST comment period on AI Agent Identity and Authorization closes on 9 March 2026. That is a formal, public, permanent record. I intend this statement to be part of it.

Whether anyone listens is not within my control. Saying it clearly and on time is.

For the Record

Patent: GB2603013.0 — Intent-Bound Authorization. Filed 5 February 2026, UK IPO.

CBID OTS: SHA256 15752d3262c58af5350bdaa2af1a7c45eaad9f636f717bd56d653e726535cc7e — 24 February 2026

NIST Submissions: NIST-2025-0035 and NCCoE AI Agent Identity & Authorization

Contact: IBA@intentbound.com | intentbound.com | governinglayer.com

+66 081 288 7843 | Chiang Mai, Thailand