Shadow AI: the risk you cannot see

Your staff adopted AI faster than your institution could govern it. The productivity is real. So is the exposure, and most of it is invisible, because the tools that create the risk are the ones no one approved. **You cannot govern what you cannot see.**

40%

of healthcare professionals have encountered an unauthorized AI tool at work. One in five admit to using them.

Wolters Kluwer, Jan 2026

$7.42M

average healthcare data breach, the costliest of any industry for fourteen years running.

IBM Cost of a Data Breach, 2025

242.9M

individuals exposed in U.S. healthcare breaches reported in 2024 alone.

HHS OCR Report to Congress

What shadow AI actually is

It is not one tool. It is every path that was never sanctioned:

- An employee pasting a record into a public chat model from a personal account.
- A coding assistant or command-line client reaching a provider that was never reviewed.
- A browser session signed into an AI tool the institution does not know exists.
- A new agent wired up by one team, invisible to everyone else.

Why most controls miss it

Most controls assume traffic takes a sanctioned path. A gateway sees only what was wired to it. A closed assistant governs only itself. Provider-side settings apply only after the data has already arrived. Each assumes away the exact traffic that defines shadow AI. The result is a control that reports clean while the real risk runs beside it, unseen.

Why it is no longer about pasted text

Modern AI tools do not just chat. They run shell commands, read and write files, query production databases, call internal APIs, and browse the web. Point one of those at a system of patient records, criminal histories, student files, or classified material, with no control in the path, and the exposure is no longer measured in a leaked paragraph. It is measured in actions taken on your most sensitive systems, with no record of who did what.

What closes the gap

Only one position sees all of it: a control on the device, at the point traffic leaves the machine, deciding before egress, across every tool and provider. Verillian sits there. It catches the sanctioned and the unsanctioned alike, decides allow, redact, or block before content leaves, and signs every event into a record you hold.

The tools you approved are not the problem. The ones you did not are. Seeing them is the whole job.