

For the CTO

An asset your team owns, with the tests attached and no lock-in.



The problem it solves

Most AI tooling decisions create a dependency: a product to renew, a prompt library nobody owns, behaviour that changes when a vendor ships. The question is what you hold in two years.

Open standard

published, not a vendor format

40+

products read it

Plain Markdown

reviewed in a pull request

How it runs

- 1 An asset, not a subscription**
A folder in the open Agent Skills format, in your repository and under your version control. It keeps working whether or not we ever speak again.
- 2 No lock-in, and reviewable in your own process**
Around forty agent products read the format, so changing editor, model or vendor costs you nothing. Plain text means a change arrives as a pull request with a diff.
- 3 Evidence, by the standard's own method**
Tested with questions your team sets, run with the skill and without it, and evaluated against the practices published alongside the specification.
- 4 Clear on data and IP**
Your material builds your skill and nothing else, is parsed in isolation, and is deleted once the work is received. No third-party extracts enter your codebase.

What you get

-  A skill folder plus the test results, with no runtime dependency on us.
-  Validation of every claim, and a mechanical check against the format's hard limits.
-  Assessment of skills you already run, scored so a later review is comparable.
-  Lean skills, loaded only where they help.

What the evidence says

- Adoption is not return: DORA 2024 found AI adoption raises individual productivity while negatively affecting delivery throughput and stability.
- Self-reported speed is unreliable: in a randomised trial developers took 19% longer with AI tools, and still believed they had been 20% faster.
- Unmeasured programmes get cancelled: Gartner expects over 40% of agentic AI projects to be cancelled by end 2027, citing cost and unclear value.

Sources: DORA, State of DevOps Report 2024 (dora.dev); METR, July 2025 (metr.org); Gartner, 25 June 2025. Context for measuring AI work, not claims about this service.