

The hidden tax on AI work isn't burnout. It's context-switching.

Every time an engineer switches tasks, they pay a tax to re-ground the AI: re-pasting context, re-explaining the goal, restarting threads that drifted. Linguist names that cost, measures it in an instrumented sandbox, and trains it down with a published rubric.

THE PROBLEM

Token spend, code review variance, "feels draining" — those are symptoms. The cause is that engineers re-pay the same context tax on every task, every interruption, every paused conversation. A team without shared patterns for re-grounding burns minutes per task, hours per week, weeks per quarter. And nobody measures it because nobody has a name for it.

WHAT WE DELIVER

LOWER CONTEXT-SWITCH COST

Headline outcome. Trained engineers re-ground faster and resume paused work cleanly.

FLOOR-RAISING

Every completer reaches the **Enabled** tier — competent, deliberate use of AI for real work.

MEASURED

Instrumented sandbox + published rubric. Bracket placement before, movement after.

METHODOLOGY

- **Free baseline session** — each engineer runs a standardised scenario in Cyborg. Output: bracket placement + org capability map + wellbeing baseline (NASA-TLX + agency/trust).
- **6-week Enabled programme** — cohort or internal-only. Tracks informed by baseline gaps. Mid-programme survey checks for drift.
- **Advanced upgrade (optional)** — for engineering teams on Claude Code: cost awareness, session hygiene, agent delegation.
- **Post-training reassessment** — same rubric, same sandbox. Report: bracket movement, behavioural deltas, wellbeing deltas.

PRICING — INDICATIVE

COMPONENT	PER ENGINEER	NOTES
Baseline audit	€0	Pre-engagement, included.
Enabled programme (6 wk)	€800	Cohort or internal-only.
Post-training assessment	€200	Same rubric, same sandbox.
Advanced upgrade	€600	Engineering teams only.

Indicative. Minimum 10 engineers. Volume terms above 50. Optional quarterly refresh available post-engagement.

OUR PROMISE

Every completer reaches the Enabled tier as defined by our published rubric, or we extend the programme at no additional cost until they do.

Linguist is the only published rubric for AI-fluency capability; Cyborg is the only instrumented sandbox built around it. **Next step:** a 30-min call to scope your team and book the free baseline — linguist.vuelolabs.com/teams