

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

Every time someone on your team picks up where they left off, they re-pay the same tax. Re-pasting context. Re-explaining the goal. Or letting one thread accumulate enough drift that quality dips and nobody catches it. Linguist names the cost. Measures it. Trains it down.

THE PROBLEM

Token spend, "feels draining", inconsistent output — those are symptoms. The cause is that everyone on your team is re-paying the same context tax. New conversation, same explanation. Or one long thread that drifts until quality dips and nobody notices when. You can feel it. You can't name it. So nobody fixes it.

WHAT WE DELIVER

LOWER CONTEXT-SWITCH COST

The headline outcome. People re-ground the AI faster and notice when a thread has gone stale.

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.

HOW IT GOES

- **Free baseline session** — each person runs a standardised scenario in Cyborg. Output: bracket placement, org-level capability map, wellbeing baseline (NASA-TLX + agency/trust).
- **6-week Enabled programme** — cohort or internal-only. Tracks informed by baseline gaps. Mid-programme survey catches drift.
- **Advanced upgrade (optional)** — for engineers using 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 SEAT	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	Engineers on Claude Code.

Indicative. Minimum 10 seats. 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