

AI Cost Commander: Granular token-level visibility into your Cortex and LLM spend

Monitor, attribute, and optimize every dollar spent on Cortex AI — before it shows up on your invoice.

THE CHALLENGE

Most teams know they're spending on Cortex AI. Few know where exactly — which model, which user, which query is driving that number up. By the time the bill arrives, spend has already compounded with no clear starting point to fix it.





OUR SOLUTION

AI Cost Commander gives you token-level visibility across every AI service, model, user, and query — with ranked optimization targets and estimated savings built in, so you act before the invoice lands, not after.

AI Cost Commander Capabilities

- 1 Credit consumption by AI service
- 2 Usage and efficiency by model
- 3 Cost attribution and risk flagging
- 4 Daily spend, volatility, and projections
- 5 cost, traffic, and adoption per agent
- 6 DAU, WAU, and Cortex Analyst usage

What Sets It Apart


-  Every dollar traced to a service, model, or individual query.
-  Built for the full Cortex stack. All AI services covered, out of the box.
-  Ranked optimization targets with estimated savings — no manual audit.
-  Know if AI is actually being used — not just deployed.

Value Delivered Across the Organization




AI & Innovation Leads

Adoption health metrics alongside spend data. Ship AI initiatives confidently, with cost trajectory fully in view.



FinOps & Finance

Full cost attribution by service, model, and user. No more end-of-month surprises or unbudgeted overruns.



Platform & Data Engineering

Query-level forensics and ranked optimization targets. Know what to fix and how much it saves before you touch anything.



CTO / VP Engineering

Executive-ready visibility into Cortex spend, agent performance, and platform utilization — all in one place.

Click Here to Watch Demo [or](#) Request a Meeting