



toknly

Every AI dollar, accounted for.

www.toknly.in

Spend intelligence for teams building on AI

A product of Pinnacle Consultancy Group

Your AI bill arrives from a dozen places. toknly turns it into one number you can act on.

OpenAI, Anthropic and Google API usage plus ChatGPT, Copilot and Cursor seats — tracked live, attributed to the team, feature and customer that caused it, and capped before an agent loop turns into an invoice.

WHAT IT DOES

Unified API + seat spend

Token spend and per-seat AI subscriptions in a single view — the only tool that reports your true, total AI cost.

Zero-risk ingestion

No proxy in your request path. Events are posted after the call, so toknly can never add latency or break production.

Per-customer economics

Tag by team, feature, customer and environment to answer "what does this customer cost to serve?" instantly.

Reasoning-token clarity

The fastest-growing, least-visible line on an AI invoice — priced and reported separately, with runaway-loop detection.

Budgets that enforce

A pre-call check returns allow / warn / block in milliseconds, so expensive calls stop before they cost you.

Fully self-serve

Sign up, pay by card or UPI, and start sending events in minutes. No demos, no sales calls, cancel anytime.

THE VALUE IT ADDS

- **Fewer surprise invoices.** Spend is visible the second a call happens, not 30 days later.
- **Correct pricing.** Set margins on real cost-per-customer instead of blended averages.
- **Overrun protection.** Budget checks cap runaway agents and misbehaving features automatically.
- **One defensible number.** Finance and engineering finally read the same AI cost report.
- **Faster model decisions.** Compare cost and reasoning-token load per model before you commit.

WHO IT'S FOR & PRICING

PLAN	PRICE	TRACKED AI SPEND
Free	Rs 0	up to \$2,000/mo
Team	Rs 6,500/mo	up to \$50,000/mo
Scale	Rs 24,900/mo	up to \$250,000/mo

- AI-native SaaS teams that price on unit economics
- Engineering leads who own model and feature cost
- Finance and ops consolidating every AI line item

LIVE IN THREE STEPS

1 • Create a key

Sign up and generate a workspace API key in Settings. Two minutes.

2 • Send events

One HTTP POST per model call with optional team, feature and customer tags.

3 • Add seats & budgets

Record AI subscriptions, set budgets — alerts and enforcement start immediately.