

PROMPTLEASH

Stop guessing your AI impact. **Start measuring it.**



Investor Pack

We're the right team

We have done this before. This team has delivered over **\$6m of work together** over the last few years across ASX Top 10 Banking institutions, designing, developing and deploying cutting edge AI platforms.



Jeremiah Mannings

Founder & CEO

Jeremiah has sold and led some of the most advanced AI platforms in Australia, personally designing and delivering over \$40m worth of AI platforms for financial institutions.

- AI Leader of the Year 2025, 3x IAPA Top 25
- Senior Director Capgemini, PwC
- Director, CDO at Uniting (\$1.1b+ rev)
- Delivered several of Australia's largest AI platforms, individually over A\$7.5m
- Founder of 6,500+ member Australian AI community



Samuel Flynn

Co-Founder & CTO

Sam has architected, developed and deployed advanced AI platforms in secure regulated environments into production as a lead software engineer.

- 8+ years enterprise software & AI system development
- High-throughput APIs, distributed infrastructure
- Lead engineer on production-ready AI platforms at scale over A\$6m



Saugata Ghosh

Chief Sales Officer

15+ yrs enterprise BD, ASX Top 10 clients, 5x revenue impact for companies like ANZ Bank, driving AI projects.

We have sold and built enterprise AI platforms before.

We know large enterprise buyers.

We have access to the right design partners.

We can actually ship.

Different enterprises, same pattern

We have spent the last 10 years in Australia's biggest banks, building platforms and enabling AI Adoption, **the same patterns have emerged**. We've repeatedly seen adoption stall within six months of go-live.



LLM provider dashboards have no indication of quality, efficiency, or compliance, just simple counts

“No idea of prompt quality, what cost is being consumed, or what people are getting caught on” - Chief Data & Analytics Officer at Westpac



Nobody can explain their AI adoption to the board, not the CTO, not the CFO

“We haven't found anything in this space, we've started looking at building ourselves” - Head of Gen AI at CommBank



No understanding if training and spend was effective, **pressure to rationalise AI spend**

“At some point we will have to deal with the cost issue, when it's no longer subsidised” - Head of Enterprise Architecture at IAG Group



Companies are reevaluating their workforces with AI, without an understanding of how it is being used



What this resulted in

68%

of licensed seats for AI tools remain unused after 90 days.

“How do I explain these costs?”

<8%

of organisations track any adoption metric beyond seat activation.

“Has AI been widely adopted?”

Zero

visibility into prompt or output quality.

“Is AI being used to its potential?”

One of our partners, ING Netherlands, is using 10 FTE for prompt quality, 5 FTE for Architecture and 5 FTE for Cyber.

Our clients are currently solving this challenge with FTE, not software.

Why is this a problem now?

AI Spend is currently subsidised

Enterprise AI infrastructure spend is projected to be **\$47B by 2027 (Gartner)**. Costs growing **3-5x YoY**. PromptLeash gives companies control ahead of the increase.

“At some point we will have to deal with the cost issue, when it’s no longer subsidised” - IAG Group

Mounting pressure to explain the outcomes

Teams largely use the basic Microsoft Copilot dashboard which shows usage, with **no idea of context and quality** of the usage.

“It would be incredible to see how each individual is using AI, and where we can improve” - Westpac

Insights beyond cost

AI will impact the workforce of the future, and PromptLeash can help organisations understand how that is progressing and where to allocate funds.

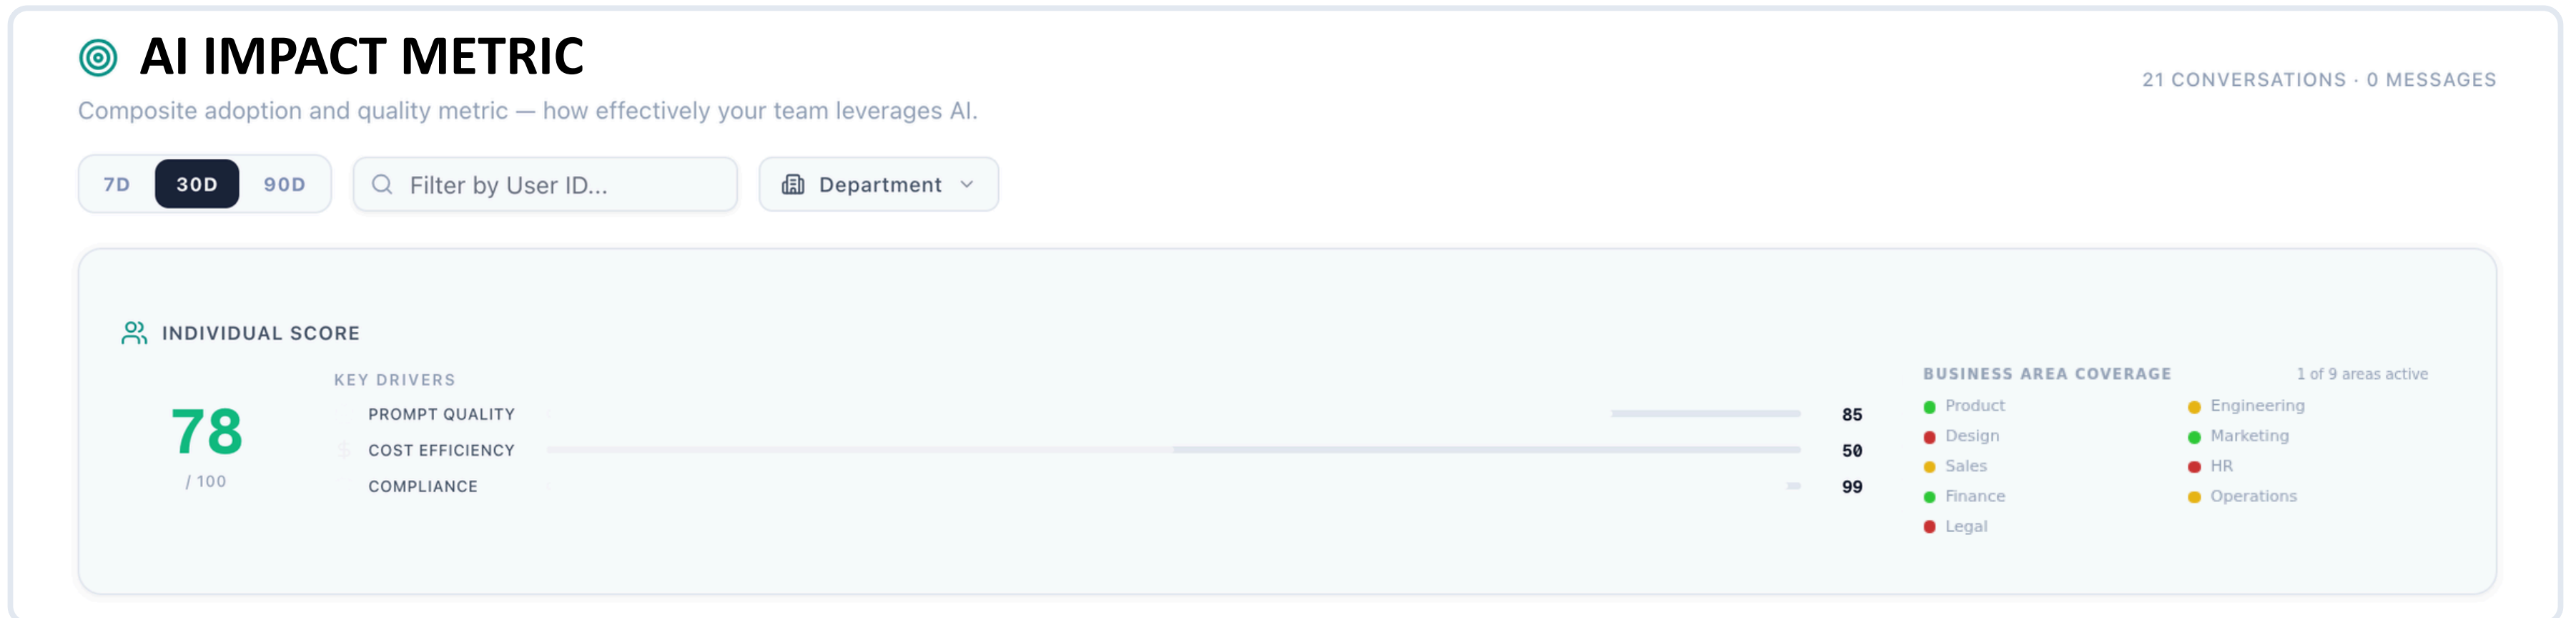
“It would be great to see a leaderboard of how our staff are using AI and creating a culture of excellence” - IAG Group

Category Validation

Portkey raised \$15M Series A. They solve technical reliability, not adoption. The AI adoption layer is **largely unowned** with no clear category leader on AI impact.

“We haven’t found anything in this space, we’ve started looking at building ourselves” - Commonwealth Bank

PromptLeash provides AI visibility



We accelerate AI adoption

Giving enterprises visibility into how AI tools are **actually being used** across their organisation, and how well it aligns to roles and culture.

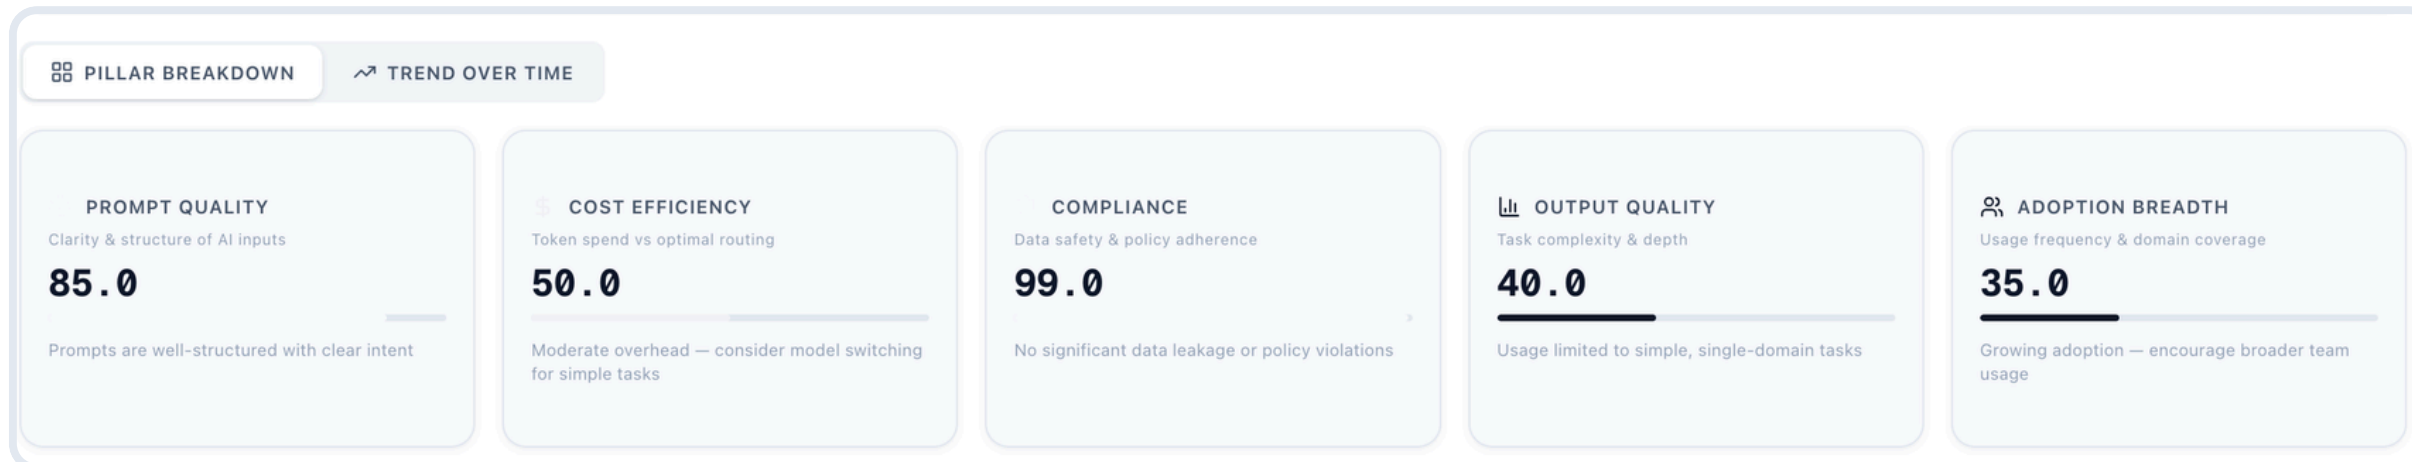
Standardised Benchmark Score

AI Impact Metric (AIM) for the enterprise, department, and individual, which shows how well individuals and teams are adopting AI.

Advanced Visibility

The AIM score allows specific analysis across different departments with **benchmarked cross-company intelligence** to similar companies.

PromptLeash shows the path forward



Score is based on three key metrics (Prompt Quality, Cost Efficiency, Compliance) with Output Quality and Adoption Breadth having minor impacts.

Immediate AIM score value

The AI Impact Metric (AIM) gives you an intuitive understanding of how well individuals and departments **score across quality, cost, and compliance**. Ultimately, how well they are doing their role with AI.

RECENT OPTIMISATIONS

▼ All Domains ▼ Most Recent ▼ Min improvement: Any ▼

Customer Complaint Response Draft

Here's a draft response: --- Dear Ms Chen, Thank you for your patience regarding the discharge of your mortgage (Reference: HL-2089-4451-7723). I sincerely apologise for the delay you have experienced. I...

+9% quality 13/04/2026, 11:28:49 ▼

Quarterly Risk Report Formatting

Here are the executive talking points for the CRO deck: **Credit Risk Position — Q1 2026** 1. **Provision Adequacy**: Total provisions of \$1.2B represent a coverage ratio of 1.82%, up 4bps QoQ, reflecting...

+8% quality 13/04/2026, 11:28:49 ▼

Actionable Recommendations

After seeing their AIM Score, organisations receive prioritised department and individual level recommendations that highlight where AI adoption is stalling and **what to do about it**.

Strong market response



In the Antler W26 Accelerator cohort

Active interest from Rampersand, Nakatomi, and active discussions with HNW investors.



Live in beta

V1 is deployed in beta, with the features to process customer logs and show AIM score and recommendations.



Enterprises building with us

IAG Group Australia (LOI signed), Westpac Group (LOI pending), Transport Accident Commission (LOI pending), CBA (LOI in progress), actively pursuing design partnerships with ING Netherlands & Admiral Group in the UK.

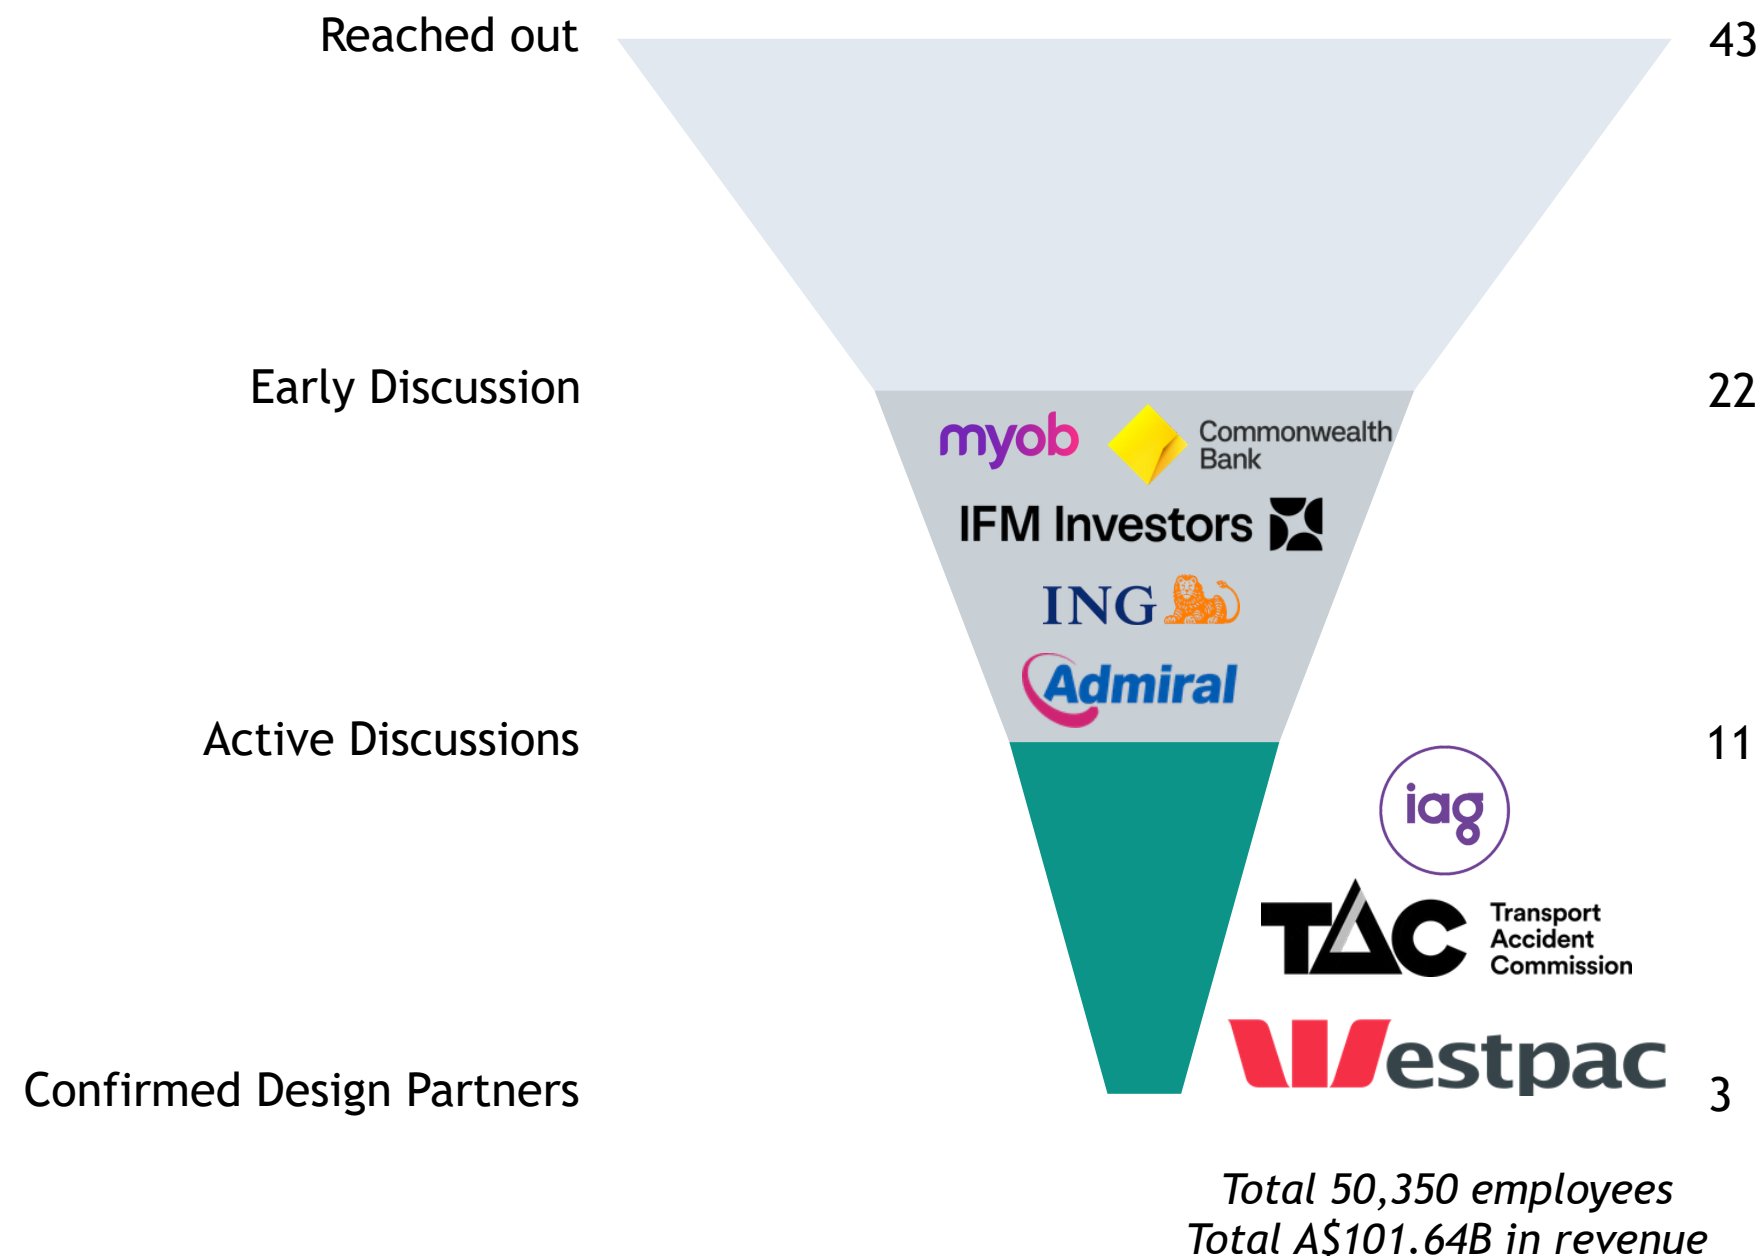


Targeted enterprise sales pipeline & GTM

Saugata (CSO) onboarded and driving pipeline and conversations active with several local and international institutions across banking, insurance and finance.

How we have validated

7 in 10 organisations are thinking about the effectiveness of their AI adoption, but have no ability to measure it.



“ Head of Agentic Transformation, ING Netherlands
(60,000 employees, A\$98.37B revenue)

“We currently have no idea of prompt quality, what is being consumed, and where users are getting caught. Our governance is fragmented, and this would help us benchmark”

“ Head of Software Engineering, Admiral Group UK
(11,000 employees, A\$11.05B revenue)

“The Engineering team is generating a vast amount of code with AI, but the real challenge is reducing the time-to-ship - there's tremendous executive pressure to demonstrate value.”

The business model

Tier 1: Partner Supported

\$20k per use

Consumed by our partner sales model, partners (consulting & technology) using our platform with their clients.

- Hosted platform on PromptLeash servers, with a tenant for specific clients
- Log upload & analysis
- Savings dashboard
- Partner led onboarding

Tier 2: Embedded Directly

\$55k platform fee

- Deployed on client cloud / environment
- PromptLeash analytics layer
- Guided onboarding (2-4 wks)
- Dedicated support
- Deployed local model in client cloud

chats per day

T1: 0-25k

T2: 25-150k

T3: 150k+

Costs per month

\$8.5k

\$15.5k

\$25.5k

Example a medium bank with 5000 people, using AI well will generate 30,000 - 35,000 chats per day. In the second tier, their costs are $\$55k + \$15.5k \times 12 = \$241k$ per year.

Pricing structure still under review

Why we win in this landscape



Speed of Category Formation

Incumbent vendors have **12-18 month product cycles**.

By the time they ship an adoption measurement layer, PromptLeash already owns the relationships, the benchmark data, and a flywheel that gets smarter with every company onboarded.

Proprietary Semantic Layer & Approach

We have the semantic layer nobody else has. Roles → What Good Looks Like → Use Cases → Patterns.

Every log we ingest improves accuracy across all customers. As we gain customers, we will expand **a uniquely valuable dataset** on AI cost-performance tradeoffs globally.

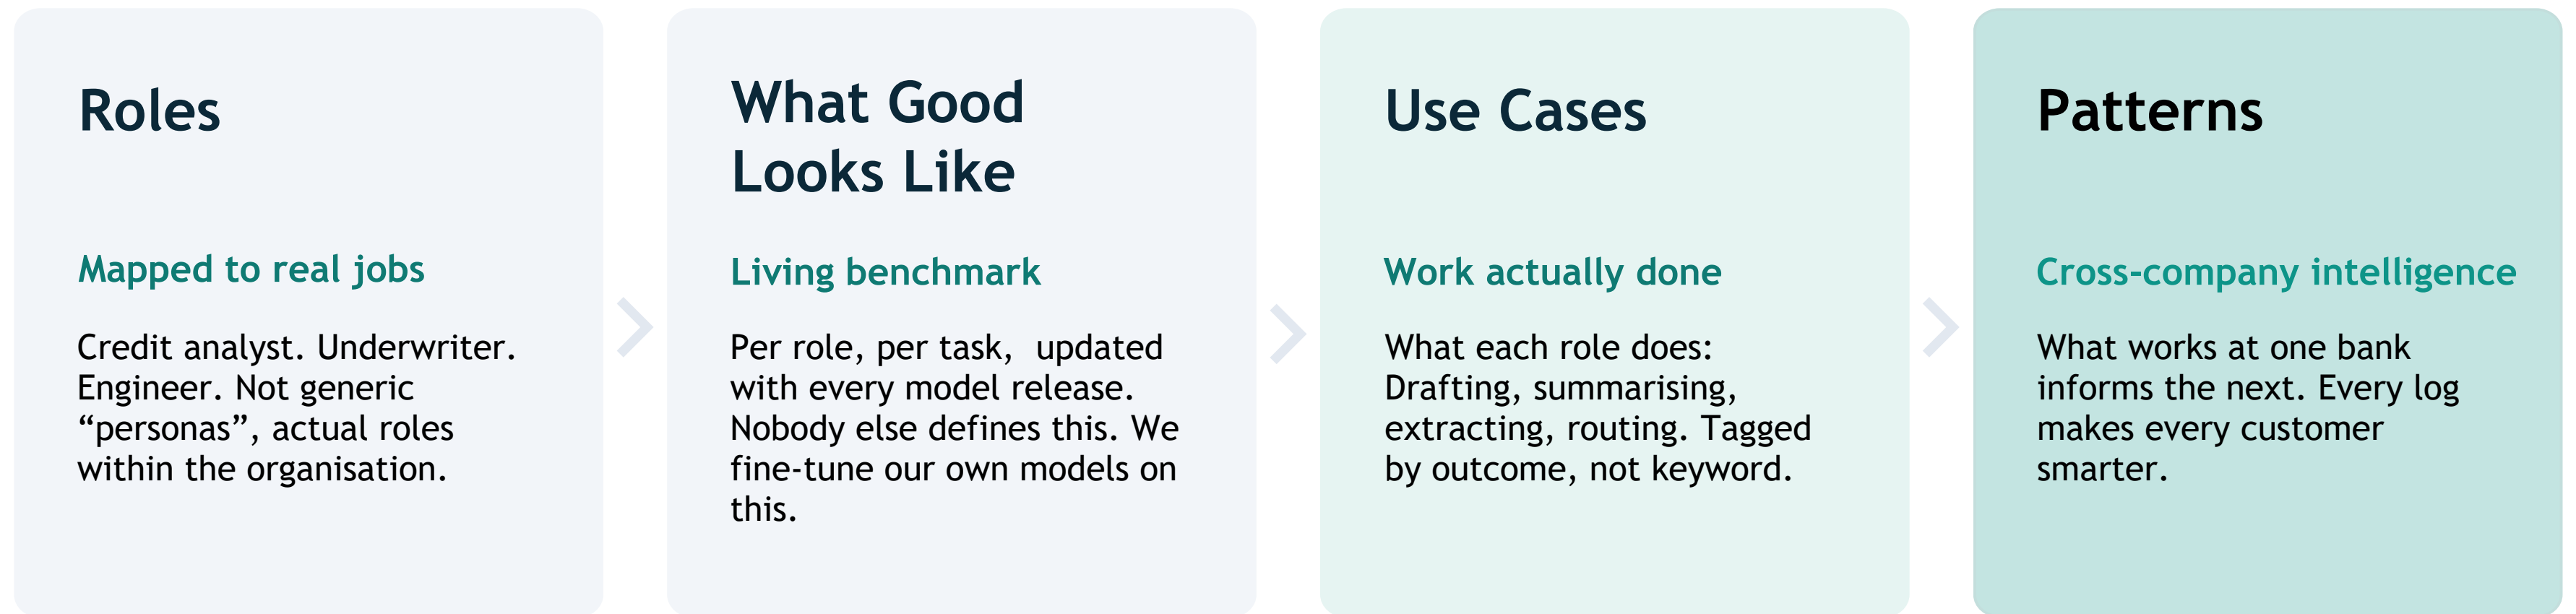
AIM Score Becomes the AI KPI

By tracking how AI is actually used across every team and workflow, **PromptLeash becomes embedded** in how the organisation measures itself.

The AI Impact Metric (AIM) Score becomes the single number that leadership points to when asking whether their **AI investment is working**.

Our primary moat

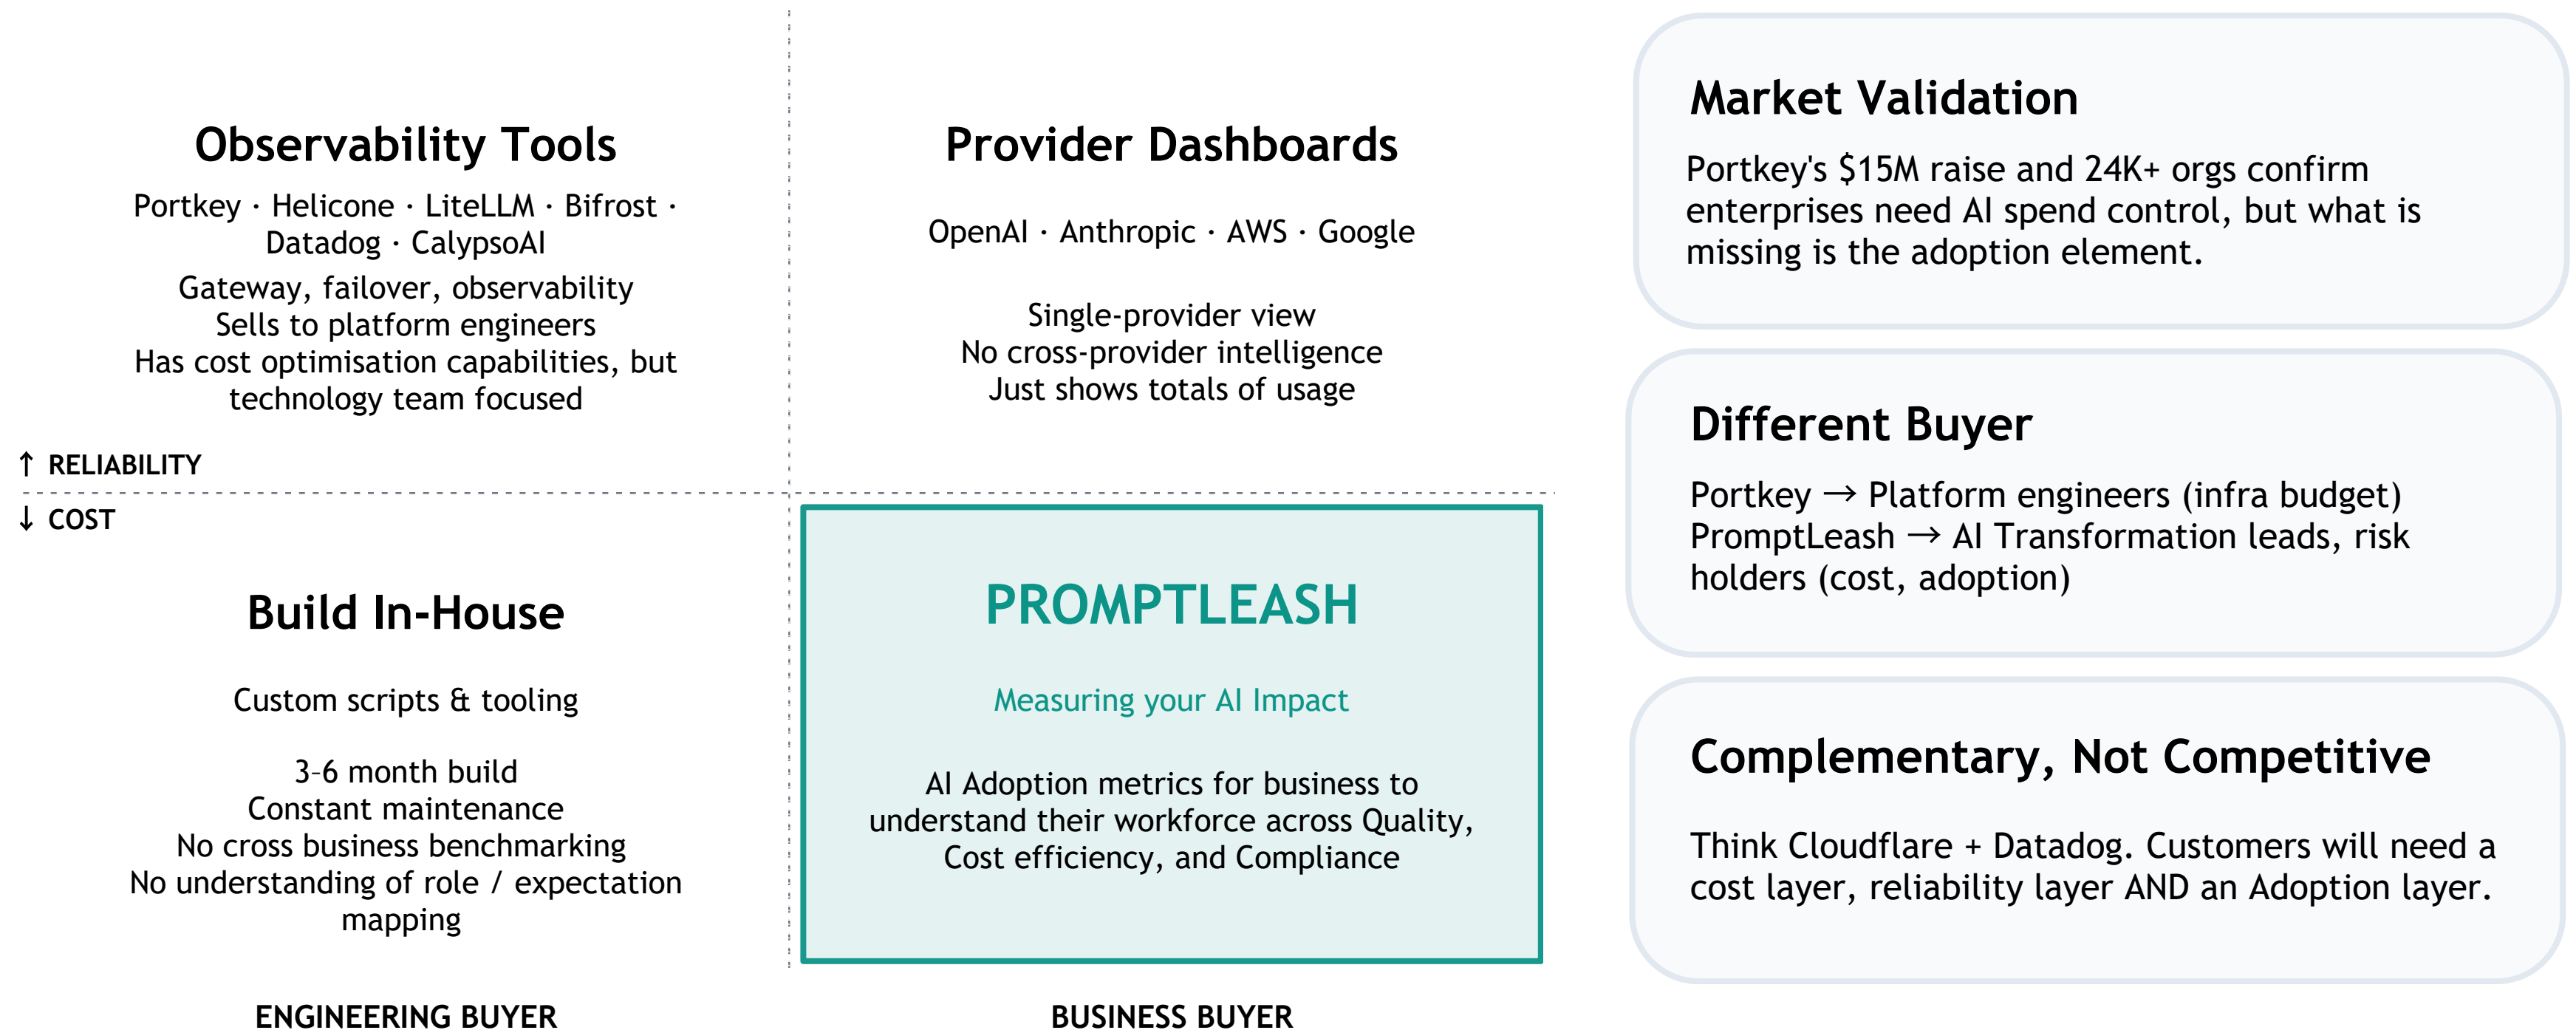
Customers keep asking what stops them building this in-house. The answer is the semantic layer.



Underpinned by Regulatory & Local Trust - Australian enterprise procurement favours local vendors with CPS 230, Privacy Act knowledge. Our Big 4 network is a structural advantage.

What good looks like is constantly changing. PromptLeash learns the benchmark across companies, roles, and workflows. Every new customer sharpens the benchmark.

Competitive landscape



Use of funds and our path forward

Our goal is to be at \$1M ARR and Series A-ready by Dec 2027.

\$2M

Seed raise

18 - 24 mo

Depending on GTM hiring pace

\$1M ARR

Series A trigger

Allocation



People 65%
GTM hires + founder salaries

Platform 15%
Infra, LLM costs, tooling

S&M 12%
Events, content, outbound

Legal/Ops 8%
Compliance, contingency

Timeline

● Apr 2026	Close design partners ✓	Founder-led.
● Aug 2026	First paid customers, IAG launch	AE + SDR onboard.
● Jun 2027	Scale to \$500K ARR	CSM hire as customers land.
● Dec 2027	Reach \$1M ARR	Expansion hires on traction.

AI optimisation has a **6-12 month horizon** where competitors will need to embed within customers and strengthen their moat. We want to do this as quickly as we can to secure this advantage.

PROMPTLEASH

The AI Adoption category is forming now.
We're building the intelligence layer.

jeremiah@promptleash.ai

promptleash.ai

End of main section, appendix follows (Risks, Architecture, ICP, Token Economics)

Risks, addressed

The four questions we get asked first. Here's how each is mitigated.

Incumbent Threat

"What stops Datadog or Portkey shipping this?"

Adoption measurement requires a semantic understanding of work, not infra metrics. Datadog doesn't know what a credit analyst does. Portkey sells to platform engineers on reliability.

By the time they ship (12-18mo cycles), we own the benchmark relationships and the cross-company dataset.

Buyer Ambiguity

"Who actually pays the invoice?"

Our wedge buyer is the CAIO, Head of AI, and AI Adoption Lead, as they own the adoption KPI. CFO is the expansion buyer once AIM Score reaches the board. CTO is the blocker-remover, not the decision-maker.

Clear sequencing: land on AI leadership budget, expand into finance ownership.

Commoditisation

"Why won't OpenAI or Anthropic offer this free?"

They structurally can't. OpenAI won't tell you to use Claude for a task. Anthropic won't route to Gemini. Neutral cross-provider measurement is a third-party job by definition.

Our dataset is cross-company and cross-provider which is a vantage point no single LLM vendor can replicate.

Geography

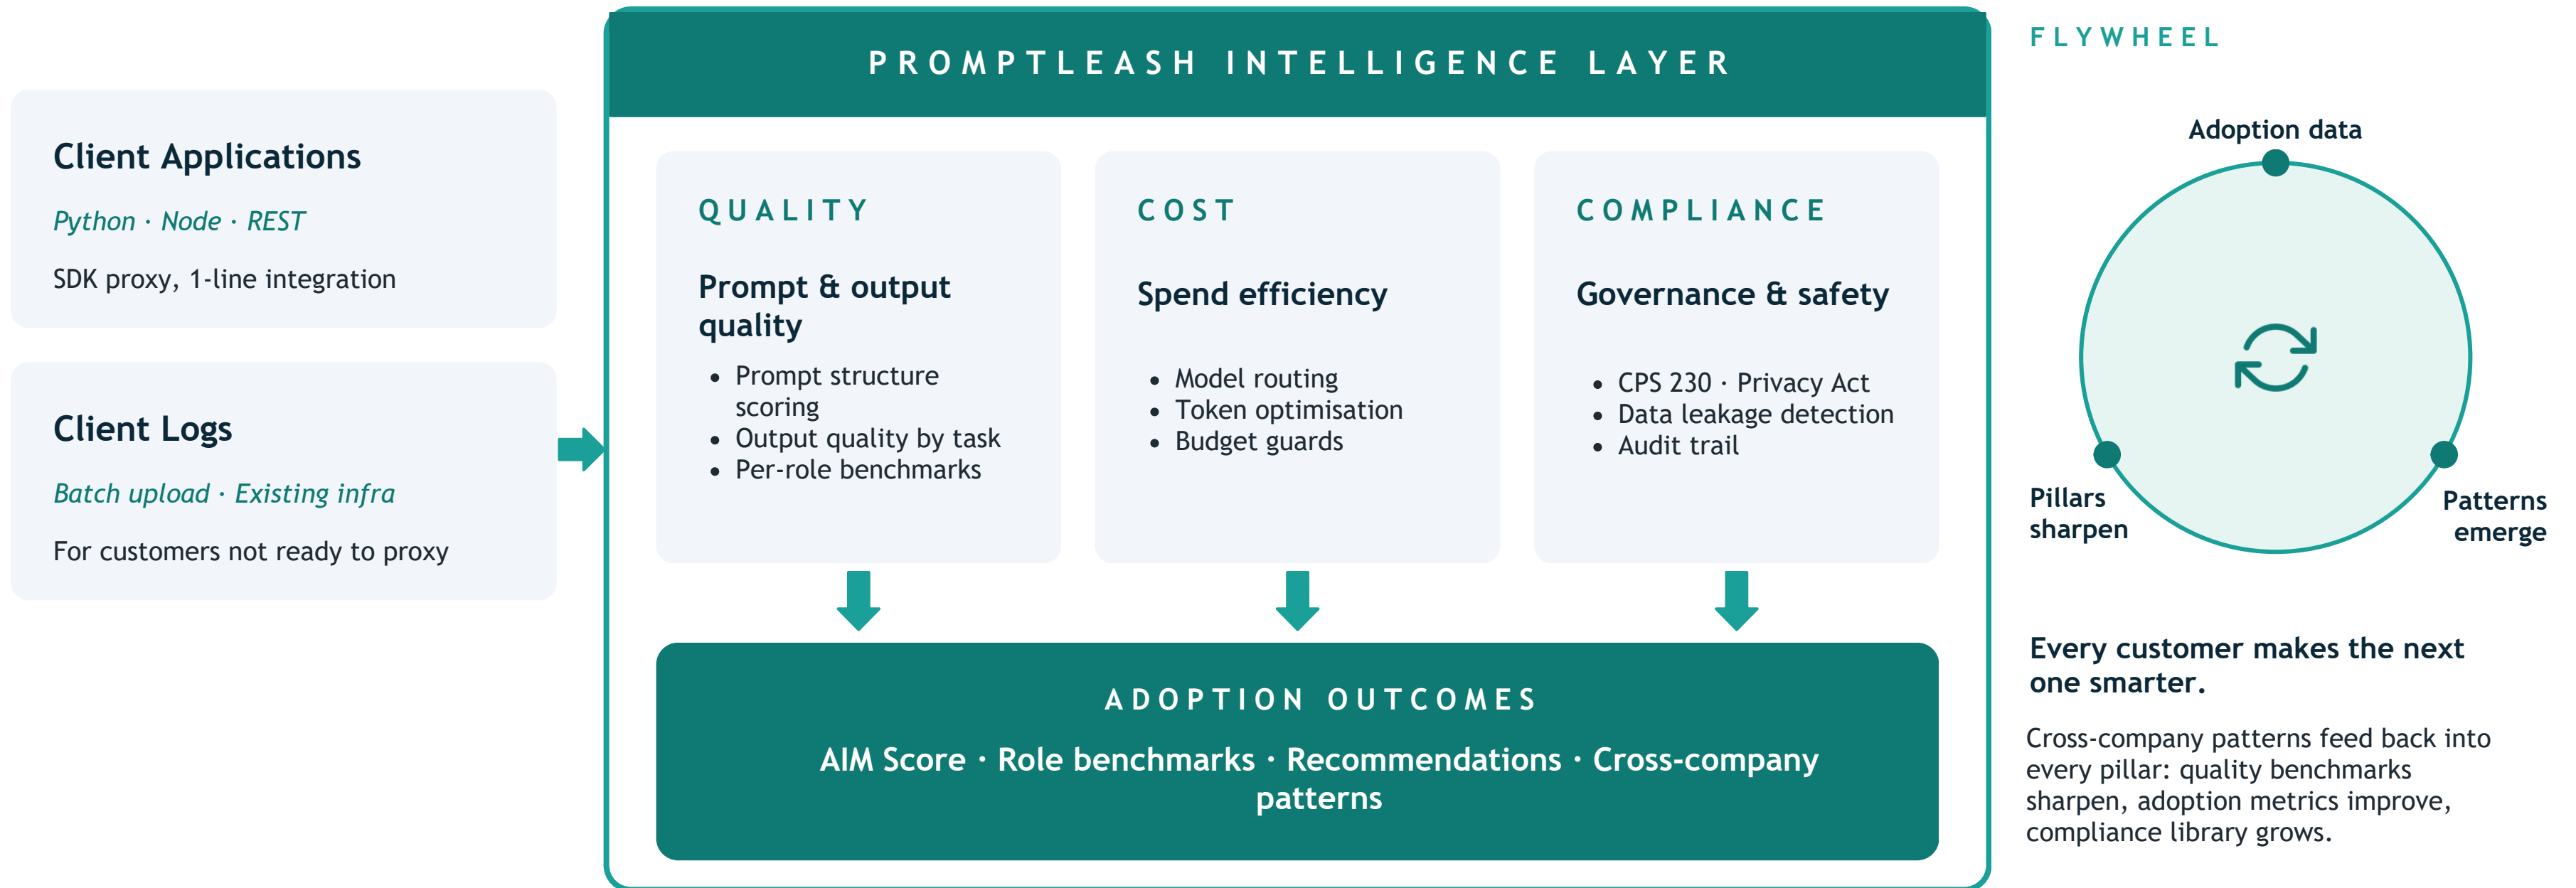
"Why start in Australia, not the US?"

CPS 230 makes AI governance forced, not optional, in APAC regulated sectors. We get to build the product with the hardest customers first, our design partners are APRA-regulated (IAG, Westpac, CBA).

ING Netherlands and Admiral UK show the product travels. We expand to softer markets with the harder product already built.

How it works

Three pillars feed one outcome. Adoption compounds the pillars. The flywheel turns.



LATENCY

Scoring runs asynchronously

DATA

No permanent storage of logs (only relevant learnings)

INTEGRATION

1 SDK change

PROVIDERS

OpenAI · Anthropic · Google · Azure · OSS

Market: ICP & TAM

~250 enterprises within our bracket today. A\$45M - A\$90M serviceable revenue opportunity in the beachhead alone.

~250

Target accounts (Wave 1)

A\$180K-360K

Target ACV, pricing scales with FTE

A\$45M-90M

Near-term serviceable beachhead

Primary: Mature Enterprise AI Adopters

Target: CAIO, Head of AI, AI Adoption Lead

Company: 1,000 - 30,000 employees, technology focused industries (FS, Insurance)

AI Spend: \$20K - \$500K/mo, growing over time

We are initially targeting APAC enterprises with sizeable AI spend, in technology focused sectors, with regulation where adoption is a key metric and focus of the executive and board.

We estimate ~250 target accounts that meet three conditions: production LLM usage across multiple teams, minimum A\$500k annual spend, central platform or risk team.

Trigger: Board AI ROI risk, rapidly growing spend, audit flags uncontrolled model proliferation.

TAM Expansion

\$1B

LLM Optimisation (2026)
(Gartner)

\$24B

AI FinOps (2027)
(Gartner / FinOps Foundation)

\$117B

Enterprise AI Management (2027+)
(IDC, McKinsey)

Token economics, the future

Token costs have a floor

Inference requires compute, memory, energy. Hardware costs rising (TSMC, NVIDIA). Hard floor driven by silicon economics.

Thinking models use 10x tokens

Chain-of-thought and reasoning models consume 5-10x more tokens per query. Per-token cost drops, total tokens per task explodes.

Enterprise AI usage compounds, subsidies reduce

Even as unit costs fall, adoption growing 3-5x YoY. More agents, more use cases, more models. Total spend rises regardless.

Our positioning strengthens

We charge on fixed tiers (adoption outcomes), not tokens processed. As complexity grows, the need for an adoption intelligence layer increases.

The Net Effect

Per-token cost



×

Tokens per task



= Total AI spend continues to rise

Adoption intelligence becomes more critical, not less.

GTM Approach

Progressive expansion in product and customer segments to get to \$1M ARR and Series A-ready by Dec 2027.

