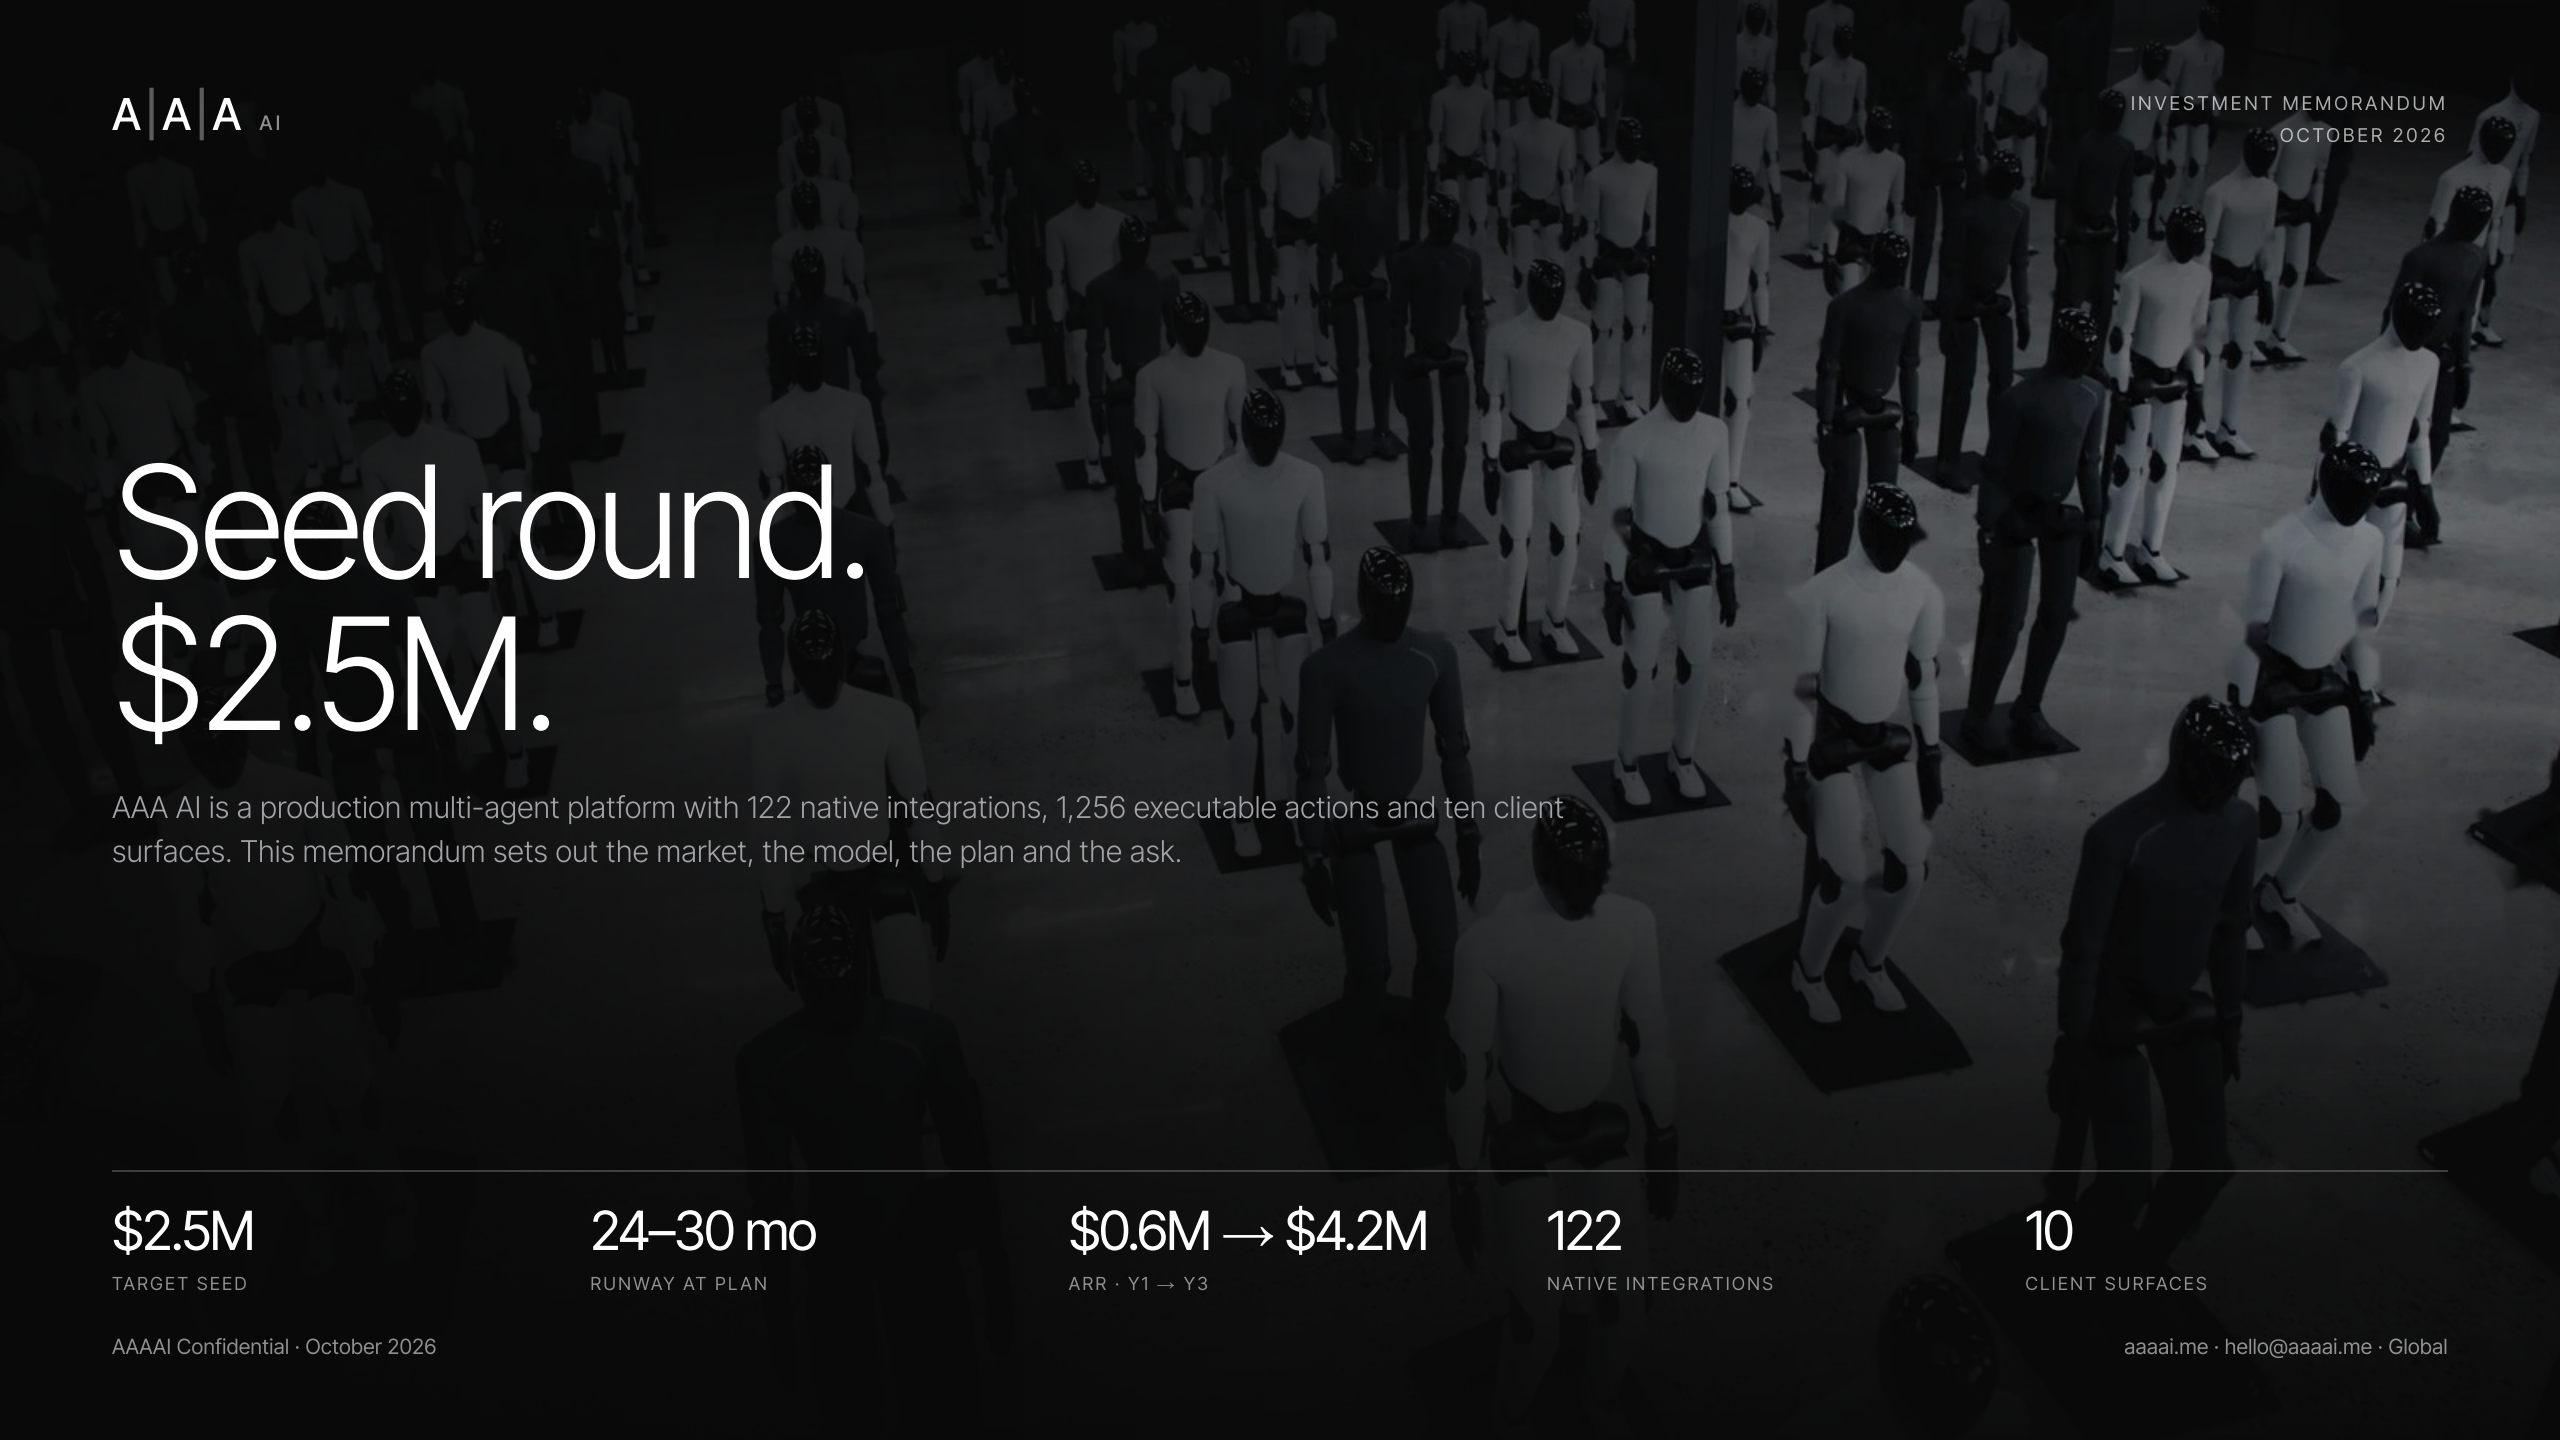


Seed round. $2.5M.

AAA AI is a production multi-agent platform with 122 native integrations, 1,256 executable actions and ten client surfaces. This memorandum sets out the market, the model, the plan and the ask.

$2.5M	24–30 mo	$0.6M → $4.2M	122	10
TARGET SEED	RUNWAY AT PLAN	ARR · Y1 → Y3	NATIVE INTEGRATIONS	CLIENT SURFACES

aaaai.me · hello@aaaai.me · Global

A shipped platform.
A revenue plan. Not a prototype.

The product is live across web, mobile, desktop, CLI and messaging. The round funds distribution and the conversion of platform breadth into repeatable Pro and enterprise revenue.

01 **What we do**

A governed workspace and runtime for multi-expert AI: chat, meetings with live notes, workflows, tasks, an AI IDE, server agents, a CLI and bots — cloud, VPC or fully offline.

02 **Where we are**

Production on AWS, ten client surfaces, public App Store listing, 122 native integrations and 1,256 actions. Monetisation is live through Paddle, Apple IAP and crypto; the AAA wallet, VISA cards (virtual live, physical soon) and the Telegram Cards app (@aaa_aimebot) are product surfaces. The current-platform audit puts replacement cost at ~$3M.

03 **Why now**

Agents can act, but enterprises cannot govern them. Approval gates, scoped credentials, memory and local execution are the missing layer — and they are already product surfaces here.

04 **The model**

Free → Pro at $20/month → Team/Enterprise. Pro is the volume engine; enterprise contracts and wallet settlement are the expansion layers. Plan: $0.6M ARR in Y1, $4.2M in Y3.

05 **The ask**

$2.5M seed for 24–30 months of runway: product 40%, acquisition 30%, infrastructure 20%, team and operations 10%.

122	1,256	$240	~10:1
Native integrations	Executable actions	Pro ARPU per year	Planning LTV / CAC on Pro

SO WHAT The risk is not whether the product exists. It is whether distribution can be funded fast enough to own the category.

Tool overload, weak AI depth, unacceptable data trade-offs.

01 / FRAGMENTATION

Nine apps, one context switch at a time

Knowledge workers move between nine applications a day. Each switch costs 23 minutes of refocus, and no assistant holds the thread between them.

02 / DEPTH

Single-model answers fail where risk lives

Math, code, compliance review and financial analysis need cross-checking. A single confident model is the wrong architecture for work that carries consequences.

03 / CONTROL

“Please don’t paste sensitive data” is not a control

Regulated buyers need local, VPC or self-hosted execution, scoped credentials, approval gates and audit-ready run history before agents touch production systems.

WHO PAYS, AND WHY NOW

- **Enterprises and regulated SMBs** — data residency, SSO, self-host and connectors into systems of record.
- **Engineering organisations** — AI IDE, agents, MCP and CLI without copying repositories into consumer chat.
- **Line-of-business teams** — repeatable research, operations and support work with a governed execution path.

The budget is already there

AI mandates are funded; what blocks deployment is trust and operations. Enterprises adopting GenAI at scale remain near 23%, while annual friction from fragmented work is estimated at $87B (McKinsey, 2025).

SO WHAT Buyers do not need another assistant. They need an operating layer that makes assistants safe to deploy.

One runtime, four layers, every surface.

LAYER 01

Surfaces

Web, iOS, Android, macOS Chat, macOS IDE, CLI TUI, desktop agents, Telegram Cards, Discord, Slack.

LAYER 02

Orchestration

Workflow builder, task board, agent teams, scheduled runs, human approvals, run history.

LAYER 03

Intelligence

10 expert modules, 13 verification modes, memory with Neptune graph, model router, search, vision, voice.

LAYER 04

Control

Ollama/MLX local execution, VPC or self-host, BYO keys, data guard, scoped credentials, audit logs.

Expert pipeline

Perception → memory retrieval → expert modules → planner → action, with cross-critique and verification before an answer is returned.

Integrations & agents

122 native integrations and 1,256 actions; a 2,048-connector catalog; hosted MCP marketplace; every action callable from chat, workflows and the CLI.

Production infrastructure

AWS with Terraform and blue-green deploys, Neptune graph, Redis, S3, usage limits, scoped API keys and billing across Paddle, Apple, crypto and card rails.

10
Expert modules

13
Verification modes

10
Client surfaces

2,048
Connectors in catalog

SO WHAT The technical risk is behind us: the remaining risk is commercial execution.

Freemium SaaS, plus agent-native settlement.

FREE

$0

Acquisition and local-first evaluation.

- Ollama / MLX models, offline
- Basic workflows and task board
- One system agent

PRO

$20 / month

Main B2C and SMB revenue driver.

- All cloud models
- 122 native integrations · 1,256 actions
- Search, image and voice
- Unlimited agents and MCP

TEAM / ENTERPRISE

$60k

Planning ACV; per-seat + installation.

- Self-host, VPC or dedicated cloud
- SSO, admin, audit trails
- Custom connectors and training
- SLA and on-prem GPU

AAA WALLET · AGENT PAYMENTS

1:1 stablecoin wallet (USDT, USDC and card top-ups), public or private settlement, and scoped agent spend keys with per-transaction and daily limits. It monetises the agent economy directly and deepens account lock-in.

REVENUE MOTION

- **PLG** free local usage → Pro conversion on cloud models and integrations.
- **Sales-assisted** enterprise pilots → security review → team rollout.
- **Expansion** seats → fine-tuning → private connectors → wallet settlement.

SO WHAT One product creates three revenue layers: seats, contracts and transaction flow.

Cards and rails add fee revenue on top of seats.

The wallet and VISA cards are live product surfaces. Card issue, load and close fees are incremental to subscriptions and are excluded from the base case until cohorts exist.

01 / WALLET

AAA stablecoin account

3% USDT/USDC and card top-ups (97% credited), 1:1 conversion, instant single-block transfers, public/private settlement and scoped agent keys.

02 / CARDS

VISA cards · virtual and physical live

Virtual $15; physical $70, issued anywhere in the world. Deposit 1%, top-up 2%, withdrawal 2%; Apple/Google/Samsung Pay, 3-D Secure, worldwide; freeze, close. Also in Telegram at **@aaa_aimebot**. Didit KYC.

03 / RAILS

Fiat, crypto and Telegram

Paddle card checkout, USDT/USDC on-ramps, card rails, the Telegram Cards Mini App (@aaa_aimebot), public explorer and PayOps reconciliation.

CARD ISSUE ECONOMICS	USD
Virtual card — issue	15.00
Physical card — planned issue	70.00
Deposit / top-up / withdrawal	1% / 2% / 2%
Close fee — ours	0.00

Virtual is live at $15 issue; physical is coming soon at $70. No opening balances. Deposit 1%, top-up 2%, withdrawal 2%; card fees are outside the base-case ARR; gas is at cost.

ILLUSTRATIVE CARD PROFIT · 10K CARDS

DRIVER	VALUE
Average annual load per card	$3,000
Top-up fee · 2%	$60 / yr
Card-fee ARR · 10k cards	~$0.60M
Contribution margin	>90%

Telegram = zero-CAC distribution

Cards-only Mini App **@aaa_aimebot** reuses the same account, KYC and card rails. No new acquisition channel or infrastructure stack.

SO WHAT The base case stands on subscriptions; cards, wallet and Telegram add high-margin fee revenue and zero-CAC distribution.

~$3M of platform, already built.

Current-platform audit, October 2026: what exists, what is live, and what it would cost to reproduce from zero. Replacement cost is an estimate, not a valuation.

AUDITED ASSET	CURRENT
Client surfaces	10
Native integrations / actions	122 / 1,256
Connector catalog	2,048
Expert modules / verification modes	10 / 13
Financial stack	Wallet · cards · Telegram
Production posture	AWS · App Store · live KYC

Scope: codebase, live surfaces, integrations, payments, KYC, cards, Telegram, deployment and tests.

REPLACEMENT COST · $M

BUILD LINE	COST
Product engineering · 14 engineer-years	2.10
Integrations & QA · 3 years	0.45
Product & design · 1.7 years	0.25
Infrastructure & security · 1.3 years	0.20
Total · ~20 engineer-years	3.00

~$3.0M
Cost to reproduce

$2.5M
Seed ask

SO WHAT The round buys distribution on top of a ~$3M replacement-cost platform, not a build from zero. Estimate only; not a valuation.

Pro pays back in a quarter. Enterprise in half a year.

Planning assumptions, to be replaced with cohort data in diligence. Pro figures use the $20/month list price and a 3.5% monthly churn assumption.

METRIC	PRO	ENTERPRISE	BASIS
ARPU / ACV	$240 / year	$60,000 / year	List price; enterprise per-seat plus installation
Gross margin	78%	80%	Local inference lowers cloud COGS; enterprise runs on dedicated capacity
Churn	3.5% / month	10% / year	Consumer-style monthly churn; annual enterprise renewal
Lifetime value	~$446	~$480,000	Gross profit × expected life
CAC	$45	$25,000	Blended paid + organic; enterprise includes solutions engineering
LTV / CAC	~9.9×	~19×	Planning target above 3×
Payback	~2.9 months	~6.3 months	Gross profit per month

Cost control

Model router shifts volume to local or cheaper providers; Pro includes fair-use limits; enterprise runs on its own infrastructure.

Expansion

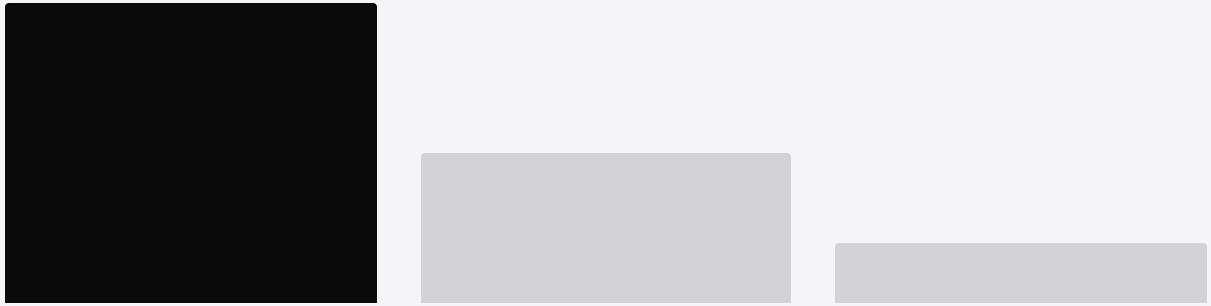
Seats, fine-tuning, private connectors and wallet settlement increase ARPU without a new acquisition cost.

Sensitivity

At 5% monthly churn, Pro LTV falls to ~$312 and LTV/CAC to ~7× — still inside a healthy band.

SO WHAT The model works at conservative churn. The round buys acquisition volume, not margin repair.

$150B+ in play.
SOM is a share we can actually serve.



TAM	SAM	SOM · Y3
$100B+	$12–18B	$4.2M ARR

TAM: global productivity software plus AI tools, order of magnitude from public industry reports. SAM: prosumer, SMB and developer teams buying privacy-aware, multi-surface AI. SOM: the Y3 plan in this memorandum.

01 **Bottom-up SOM**

2,000 Pro + 2 enterprise in Y1 → 13,000 Pro + 18 enterprise in Y3. At list price this is $4.2M ARR — 0.03% of SAM.

02 **Why the share is credible**

Distribution is already built: app stores, direct downloads, CLI, bots and a site designed for AI-agent discovery. The constraint is conversion spend, not reach.

03 **What would expand it**

Enterprise references, SOC 2 readiness, regional payment rails and connector depth toward 500+ native integrations.

400M+
Knowledge workers globally

~23%
Enterprise GenAI adoption at scale

SO WHAT A small share of a very large market is enough: 0.03% of SAM funds the next stage.

Three years to $4.2M ARR and breakeven EBITDA.

Planning scenario, not a forecast. The model assumes 10% free→Pro conversion at scale, a $60k enterprise ACV and a 3.5% monthly Pro churn.

$ MILLIONS	Y1	Y2	Y3
Ending ARR	0.60	1.80	4.20
Revenue	0.42	1.25	3.10
Gross margin	70%	76%	80%
Gross profit	0.29	0.95	2.48
Operating expenses	1.10	1.85	2.65
EBITDA	(0.81)	(0.90)	(0.17)
Headcount (end)	4	7	10
Paying Pro (end)	2,000	6,000	13,000
Enterprise (end)	2	6	18
Net new ARR	0.60	1.20	2.40

ENDING ARR · $M



3.0×
ARR growth Y2 → Y3

(1.88)
Cumulative EBITDA burn · $M

What must be true

10% free→Pro conversion at scale, a $60k enterprise ACV, 3.5% monthly Pro churn and gross margin expanding from 70% to 80% as local inference absorbs volume.

SO WHAT Cumulative burn of $1.88M leaves a buffer inside the $2.5M round and a path to breakeven in Y3.

Seats build the base. Contracts build the value.

DRIVER	Y1	Y2	Y3
Paying Pro (end)	2,000	6,000	13,000
Pro ARR	$0.48M	$1.44M	$3.12M
Enterprise contracts	2	6	18
Enterprise ARR	$0.12M	$0.36M	$1.08M
Total ARR	**$0.60M**	**$1.80M**	**$4.20M**

Pro ARR = paying seats × $240. Enterprise ACV held flat at $60k; price expansion is upside, not base case. At 10k cards, illustrative card-fee ARR is ~$0.60M at >90% gross margin, excluded from the plan.

ILLUSTRATIVE PLG FUNNEL AT SCALE



What is not in the base case

Wallet top-up fees, card issue/load/close fees, Telegram Cards distribution, marketplace take rates, training/GPU add-ons and reseller margin. Each is a separate upside line, not a dependency of the plan.

SO WHAT Seventy-four percent of Y3 ARR is Pro volume; enterprise is the margin and credibility layer.

$2.5M, allocated to distribution and delivery.



LINE	$M	SHARE
Product & engineering	1.00	40%
Acquisition & GTM	0.75	30%
Infrastructure	0.50	20%
Team & operations	0.25	10%
Total	2.50	100%

01 **Product**
Founding engineer and ML/evaluation engineer; connector depth toward 500+ native; enterprise controls.

02 **Acquisition**
App-store and search growth, founder-led enterprise, partner and channel programmes, bilingual EN/RU.

03 **Infrastructure**
GPU capacity for local-first and enterprise pilots; observability, reliability and security reviews.

04 **Team & operations**
Support, finance, legal and the compliance artefacts enterprise buyers require.

24–30 mo
Runway at plan

2 → 10
Headcount

18
Enterprise contracts by Y3

500+
Native connector target

SO WHAT Every line maps to a measurable output: seats, contracts or connector depth.

Product-led in. Enterprise-led up.

Four quarters from raise to a repeatable enterprise motion, with PLG carrying volume and founder-led sales carrying contract value.

Q4 2026

Close and instrument

Round close, hiring, funnel instrumentation, investor materials and data room.

Q1 2027

Pro engine

App-store and search campaigns; onboarding and activation work; 1,000 paying Pro.

Q2 2027

First enterprise pilots

Two design partners in regulated or mobile-heavy roles; security review pack.

Q3 2027

Team rollout

First paid enterprise contracts; SSO, audit exports and admin controls hardened.

Q4 2027

Repeatability

2,000 Pro, 2 enterprise contracts, $0.6M ARR; documented playbook for Y2.

CHANNELS

App Store / Play · Direct downloads · AI-native discovery · Founder-led outbound · Telegram Cards · Partners & resellers · EN + RU

Sales motion

Land with a scoped workflow and a defined data boundary; expand seats and connectors; convert to annual enterprise with SLA and support.

SO WHAT The plan does not depend on a single channel: app stores, search, outbound and partners run in parallel.

What “on plan” looks like every quarter.

QUARTER	PAYING PRO	ENTERPRISE	ARR	KEY MILESTONE
Q4 2026	600	—	$0.14M	Round closed; funnel instrumented
Q1 2027	1,000	—	$0.24M	Activation and onboarding shipped
Q2 2027	1,400	1 pilot	$0.40M	Security review pack complete
Q3 2027	1,700	1 paid	$0.47M	First enterprise contract signed
Q4 2027	2,000	2	$0.60M	Y1 plan; playbook documented
Y2 exit	6,000	6	$1.80M	Repeatable mid-market motion
Y3 exit	**13,000**	**18**	**$4.20M**	**EBITDA breakeven trajectory**

Leading indicators

Activation rate, time-to-first-workflow, connector installs per account and weekly active seats.

Guardrails

Model cost as a share of Pro revenue, support tickets per 100 seats and gross margin by cohort.

Diligence artefacts

Monthly plan vs actual, cohort retention, ARR bridge and infrastructure cost per active user.

SO WHAT The model is managed quarterly: seats, contracts and ARR are reported together, not in isolation.

Three categories, one combined position.

CAPABILITY	CHATGPT / CLAUDE	MS COPILOT	N8N / ZAPIER	AAA AI
Multi-expert verification	Single model	Single model	No AI core	10 experts · 13 modes
Local / offline execution	Cloud-first	Cloud-first	Self-host runner	Ollama · MLX · llama.cpp
Native connectors	Plugins	M365 Graph	500+ apps	122 native · 1,256 actions
Workflows + approvals	Agent beta	Copilot Studio	Strong automation	Nodes + human gates
Surfaces	Apps	Office	Web-first	10 surfaces
Agent payments	No	No	No	Wallet · spend keys
Self-hosted	SaaS only	Azure-bound	n8n self-host	Docker · Terraform · AWS

Breadth

Each connector and surface makes the next workflow cheaper to deploy and harder to migrate.

Governance

Approvals, scoped keys, spend limits and run history are product surfaces, not a compliance appendix.

Switching costs

Memory, workflows, credentials, team lanes and wallet balances accumulate inside the account.

SO WHAT No competitor combines an AI core, automation breadth and ten distribution surfaces.

Founder-led delivery. Hiring against the plan.

A|A|A AI

AAA AI

FOUNDER · ARCHITECT · IMPLEMENTER

More than 15 years at VP and Chief Cloud Architect level across AWS, Azure and GCP, leading teams of 15–50+ on venture-backed platform engagements. Lead architect and hands-on implementer of the full AAA AI stack: model routing, expert pipeline, memory, workflows, native clients, CLI and billing.

15+	50+	10
Years at VP / Chief Cloud Architect level	Largest organisation led	Surfaces shipped

HIRING PLAN WITH THE ROUND

01 **Founding engineer, platform**
Orchestration, connectors, reliability.

02 **ML / evaluation engineer**
Expert routing, verification quality, evals.

03 **GTM lead**
PLG funnel, app-store growth, partnerships.

04 **Enterprise solutions architect**
Pilots, security review, deployment.

05 **Support engineer**
Pro and enterprise response, docs.

SO WHAT The founder has operated enterprise infrastructure at scale and built this platform end to end.

What could go wrong. What we do about it.

RISK	WHY IT MATTERS	MITIGATION
Model cost volatility	Gross margin on Pro depends on inference prices.	Multi-provider router, local fallback, fair-use limits and per-cohort margin tracking.
Enterprise sales cycle	Longer cycles delay the contract layer.	Design partners, security pack prepared upfront, staged pilots with success metrics.
Platform competition	Large vendors can bundle assistants and automation.	Win on governance, local execution and ten surfaces; stay model-agnostic.
Key-person dependency	The founder is the principal architect.	Founding engineer hire, documentation, AI-augmented delivery and contractor bench.
Security incident	Trust is the product's core asset.	Data guard, scoped credentials, least privilege, audit logs, independent review before enterprise scale.
Payment concentration	Regional rails can be disrupted.	Paddle, Apple, crypto and card rails run in parallel; Telegram Cards adds zero-CAC distribution and wallet settlement diversifies further.

This is a planning document. Financial figures are illustrative scenarios, not forecasts, and must be reconciled with the term sheet, cap table and data room before external distribution.

SO WHAT The principal risks are commercial and operational — both are funded and managed by the plan.

$2.5M seed.
24–30 months of runway.

$2.5M

Target round · terms in the data room

- **Use of funds** product 40% · acquisition 30% · infrastructure 20% · team and operations 10%.
- **Plan** $0.6M ARR in Y1, $1.8M in Y2, $4.2M in Y3; EBITDA breakeven trajectory in Y3.
- **Milestones** 2,000 paying Pro and 2 enterprise contracts by the end of Y1.

DATA ROOM INDEX

- Cap table, term sheet and incorporation documents
- Three-year model with assumptions and sensitivity
- Architecture, security one-pager and sub-processor list
- Product demo script and live environment access
- Connector catalog export and release history

aaaai.me · web.aaaai.me · hello@aaaai.me

"The platform is already live. The round decides how fast it reaches the market."

INVESTMENT SUMMARY · OCTOBER 2026

1 · Demo

45 minutes on a real workflow with your data boundary.

2 · Diligence

Architecture, security and model inventory review.

3 · Close

Term sheet, data room and 30-day close.

SO WHAT The ask funds a measurable 12-month plan on a platform that is already live.