



WARBLE CLOUD

SOVEREIGN AI INFRASTRUCTURE


Warble Cloud

Your OWN Cloud and AI Tooling.
Zero lock-in. Total autonomy.

Pitch Deck | October 2026

Avinash Patil, Founder

warblecloud.ai • warblecloud.com • frakma.io



Rented AI. Locked-in cloud. Slow recovery.

Lock-in

AI wrappers and managed services own the data path, the models, and the roadmap. Switching costs compound every quarter.

87%

of enterprise LLM workloads run on closed-source models

MENLO VENTURES, JUL 2025

Cloud waste

Kubernetes and GPU fleets get sized for peak and billed for idle. At scale, that waste becomes a real line item.

27%

of IaaS and PaaS cloud spend is estimated as wasted

FLEXERA, STATE OF THE CLOUD 2025

Slow MTTR

Incidents still mean humans grepping logs at 3 AM. Recovery is measured in hours, not minutes.

\$2M/hr

median cost of a high-impact IT outage

NEW RELIC, OBSERVABILITY 2025

Result: teams rent their intelligence layer and pay for it twice, in fees and in downtime.

The ownership window is open

01

Models you can own

Open models are production-grade. Running inference in your own account is now a real option, not a research project.

02

MCP is the interface

Model Context Protocol is becoming the standard way agents operate tools. Infra that speaks MCP gets automated first.

03

GPU spend is board-level

GPU capacity is the most expensive line in the AI budget. Utilization is now a CFO question.

04

Sovereignty pressure

Data residency and security reviews push AI workloads back inside the customer VPC.

\$80B

sovereign cloud IaaS spend forecast for 2026, up 35.6% year over year

GARTNER, WORLDWIDE SOVEREIGN CLOUD IAAS FORECAST, FEB 2026

One platform. Your VPC. Your models.

Murmuration fixed-price architecture advisory (entry point)

Starling

MCP server for
Kubernetes ops · CLI

Reflexion

Actor/Critic incident
fix, shadow mode first

Army of Agents

Pay per use on your
cloud

Your account · VPC · Kubernetes · GPUs

- **Deploys in your account**
Infra, telemetry, and models stay yours.
- **Zero lock-in**
OSS-first. Walk away with everything.
- **Agent-ready by default**
Every capability exposed via MCP and CLI.

Advice → credits → owned platform

What Warble solves, problem first

THE CUSTOMER PROBLEM

WHAT WARBLE DOES

Paged at 3 AM, then hours of grepping logs to recover.

→ **Reflexion**

Actor/Critic agents diagnose and draft the fix. Shadow mode is on by default.

Simple cluster questions need kubectl skills and five dashboards.

→ **Starling**

Ask your cluster in plain words through the Starling MCP server or CLI.

Unsure which stack to choose, and tech debt keeps piling up.

→ **Murmuration**

Fixed-price advisory that ends in a written plan, with credits included.

ML pipelines stitched together by hand, hard to reproduce.

→ **FRAKMA**

Self-hosted, GitOps-native MLOps on Kubernetes. Azure and GCP ready; AWS in testing.

Ops agents that only run in someone else's cloud.

→ **Army of Agents**

Pay per use agents on your own cloud, at GA after November 2026.

Four building blocks, one control plane

KUBERNETES OPS VIA MCP

STABLE VERSION READY

Starling

- MCP server and CLI: agents operate Kubernetes safely
- Starling-EX beta tested; stable version ready
- External features and integrations in development

AGENTIC INCIDENT RESPONSE

CLOSED BETA

Reflexion

- Actor/Critic diagnosis and fix
- Shadow mode on by default: proposes, does not apply
- Critic checks blast radius before any action

OPS AGENTS

GA AFTER NOV 2026

Army of Agents

- Specialized agents that observe, reason, and act
- Free to start, pay per use on your cloud
- Unlocks at GA, after November 2026

AIOPS MODEL FAMILY

NOT TRAINED YET

Warble Brain (Avirka)

- Warble's own SRE models for the Reflexion stack
- Designed for the FRAKMA and Reflexion stack
- Not trained yet

Murmuration is the front door



Why it works: advisory revenue pays for acquisition, and bundled credits turn a recommendation into platform usage on day one. Buyers start with a fixed price, not a sales cycle.

Engineering that moves like a murmuration

For teams unsure which stack to choose and carrying tech debt. Fixed-price architecture advisory with written docs and included Starling credits. It is the gateway into Warble products.

FOUNDERS AND SOLO ENGINEERS

Niche Advisory

\$500

5,000 credits

One deep-dive session and a written recommendations doc

SEED TO SERIES A TEAMS

Growth Engagement

\$2,500

25,000 credits

Three architecture sessions and a build plan you can ship

SERIES B+ ENGINEERING ORGS

Scale Partner

\$7,500

75,000 credits

Full architecture design and build plan, Reflexion pilot setup

MID-MARKET, REGULATED

Enterprise


Talk to us

Custom

Custom contract, SLA-backed, embedded with your team

Q4 Growth Partner slots are open: 25 Startup and 25 Enterprise seats. Startup tier from **[to fill: Startup tier price]**.

Actor proposes. Critic verifies. You own it.



YOUR VPC

- Agents, telemetry, and models run in your account
- Starling MCP gives agents scoped, auditable access
- Shadow mode is on by default: fixes are proposed, not applied
- Every fix is a reviewable change, not a black box

Every action is policy-gated and reversible

Pre-revenue, building in the open

\$15K

raised to date, invested in compliance

Pre-revenue

in validation phase with Linux Foundation and AALF communities

Beta

Starling-EX beta tested; stable version ready

Product status

- Reflexion: closed beta, shadow mode on by default
- FRAKMA: Azure and GCP ready; AWS in testing
- Murmuration checkout live in Stripe test mode
- SOC 2: in progress

Company

- warblecloud, LLC formed via Stripe Atlas (Delaware)
- Peachscore Cohort
- Silver Member: Linux Foundation and AALF
- Design partners / pilots: **[to fill]**

Membership is not a certification of Warble products.

Fixed-price entry. Usage-based expansion.

Tier	Price	Credits granted	Top-up value
Niche Advisory	\$500 one-time	5,000	\$100
Growth Engagement ★	\$2,500 one-time	25,000	\$500
Scale Partner	\$7,500 one-time	75,000	\$1,500
Enterprise	Talk to us	Custom	Custom

\$1 = 50

Credit top-up rate. Murmuration tiers grant a flat credit amount; top-ups add more at \$1 = 50.

★ *Recommended tier*

Revenue layers: advisory (one-time) → credit top-ups and seats (recurring) → enterprise contracts. Army of Agents: pay per use on customer cloud, GA after November 2026. Seat pricing: Startup tier from **[to fill: Startup tier price]**.

Own it, don't rent it

	Warble	Managed AI wrapper platforms	Hyperscaler managed AI services	Traditional AIOps / observability
Where it runs	Your account, VPC and Kubernetes	Vendor's cloud	Your cloud account, provider-run control plane	Mostly vendor SaaS; collectors in your env
Lock-in	OSS-first. Walk away with everything	Prompts, data and workflows live in their platform	Tied to one cloud's APIs and model catalog	Telemetry, dashboards and alerts live in their platform
Agent interface	MCP-native: Starling MCP server and CLI	Vendor API or plugins; MCP support varies	Provider agent frameworks; MCP support varies	Built-in AI assistants; MCP support varies
Incident response	Actor/Critic agents, shadow mode by default	Not the focus	Building blocks; you assemble the loop	Detect and correlate; fixes stay manual or runbook
Entry price	Fixed-price advisory from \$500	Per seat or per token	Pay as you go at provider rates	Per host, data volume or seat
Honest trade-off	Pre-revenue; you operate it in your cloud	Fastest path to a first demo	Scale, compliance, one-bill procurement	Mature, broad telemetry coverage

Category-level comparison of typical offerings; individual vendors vary. MCP = Model Context Protocol.

Infrastructure builders, led by the founder



Avinash R Patil

FOUNDER AND CHIEF FLOW OFFICER

12+ years in cloud-native platforms and DevOps. CKA. Builder of frakma.io. Agentic LLM Ops, FinOps, shift-left perimeter security; Warble Platform for AI, Reflexion Engine, Starling-EX.



Rasika Khandare

PRINCIPAL CONSULTANT, WARBLE CLOUD · DESIGNATED PARTNER, CHIRPSTACK LLP

Collaboration, business and release management, GTM, and Salesforce integration development. Azure expert for frakma.io.



Khushi Khetan

AI / DATA ENGINEER

LLM orchestration. Reflexion reasoning core.



Yaksha Tikoria

MLOPS / LLMOPS ENGINEER INTERN

MLOps infra: model serving, vector DBs, observability. CI/CD with Jenkins, GitHub Actions and Docker.



Akshat Dwivedi

AI / ML ENGINEER INTERN

GenAI and agentic AI: LLMs, RAG, LangGraph, MCP, FastAPI. Prior: GenAI intern at OxiqAI, backend intern at Xebia.

Advisors and partners around the core team



Prashant Nakhate

SOLUTION ADVISOR

Observability and SRE, multi-cluster monitoring.



Ankur Jain

SOLUTION ADVISOR

GitOps and platform engineering, Argo CD.



Umang Atal

UX DESIGN

Product design for AI ops dashboards.



Smit Mehta

SOLUTION ARCHITECT

AI stack architecture, token-efficient infrastructure.



Ankit Panwar

[to confirm]

0-to-1 product builder. Ex co-founder, COO and CTO, De Nada Tecnologia.

HOW THE COMPANY IS SET UP

warblecloud, LLC: Delaware company, formed via Stripe Atlas.

Engineering in India through Chirpstack LLP.

Silver Member of the Linux Foundation and the Agentic AI Foundation (AAIF). Peachescore Cohort.

Open source roots. Foundation member.

MEMBERSHIPS AND PROGRAMS

SILVER MEMBER

Linux Foundation

SILVER MEMBER

Agentic AI Foundation (AAIF)

FORMED VIA

Stripe Atlas (Delaware)

ACCELERATOR

Peachscore Cohort

warblecloud, LLC is a Silver Member of the Linux Foundation and the Agentic AI Foundation, in validation phase with both communities. Membership is not a certification of Warble products.

WEB AND OPEN SOURCE

warblecloud.ai

Main site: Reflexion, the agent flock, Fitment calculator

warblecloud.com

Murmuration advisory, Fitment quiz, product pages

frakma.io

FRAKMA: self-hosted, GitOps-native MLOps on Kubernetes. Azure and GCP ready; AWS in testing.

starling-ex.frakma.io

Starling-EX Community Edition: Apache-2.0 Kubernetes platform with operator, Helm chart, and Dex SSO

Start with one pilot

DESIGN PARTNERS

1. Take the Fitment quiz at warblecloud.com
2. Book Growth Engagement: \$2,500, 25,000 credits
3. Run Reflexion beta in shadow mode on one service
4. Measure MTTR and cloud spend before and after

INVESTORS

- Raised to date: \$15,000, invested in compliance
- Next round: **[to fill: target amount]**
- Intros to platform teams running Kubernetes and GPUs
- Partners who back OSS-first, anti lock-in infra



WARBLE CLOUD

Stop Renting AI Wrappers. Own Your Cognitive Infrastructure.

Avinash Patil

Founder and Chief Flow Officer, Warble Cloud

avinash@warblecloud.ai · avinash@frakma.io · sales@warblecloud.com

warblecloud.ai · warblecloud.com
frakma.io · starling-ex.frakma.io


SILVER MEMBER Linux Foundation

SILVER MEMBER Agentic AI Foundation (AAIF)

FORMED VIA Stripe Atlas

ACCELERATOR Peachscore Cohort

warblecloud, LLC (Delaware) · Engineering in India: Chirpstack LLP



Where the numbers come from

87%

SLIDE 2

Menlo Ventures, 2025 Mid-Year LLM Market Update (Jul 2025). Share of enterprise LLM production usage.

menlovc.com/perspective/2025-mid-year-llm-market-update

27%

SLIDE 2

Flexera, 2025 State of the Cloud Report (Mar 2025). Self-estimated wasted IaaS and PaaS spend.

flexera.com/about-us/press-center/new-flexera-report-finds-84-percent-of-organizations-struggle-to-manage-cloud-spend

\$2M/hr

SLIDE 2

New Relic, 2025 Observability Forecast (Sep 2025). Median cost of a high-impact outage.

newrelic.com/press-release/20250917

\$80B

SLIDE 3

Gartner, Worldwide Sovereign Cloud IaaS Spending Will Total \$80 Billion in 2026 (Feb 2026). Forecast.

gartner.com/en/newsroom/press-releases/2026-02-09-gartner-says-worldwide-sovereign-cloud-iaas-spending-will-total-us-dollars-80-billion-in-2026

All four figures are cited in Warble's Industry Problem Statement deck (Oct 2026, 18 sources). Warble does not publish its own customer metrics yet: pre-revenue.