

MURMURATION · ARCHITECTURE ADVISORY

A fixed-price read on your architecture, with a written plan at the end.

For teams unsure which stack to choose and carrying tech debt. Every tier has a fixed price, ends in a written document, and includes Starling credits so your team can keep working between sessions.

Tiers approved

Checkout: Stripe test mode

THE PROBLEM

You know the stack needs to change, but every option comes with a sales team. Freelancers bill by the hour; big integrators start with a multi-month SOW. Meanwhile, tech debt compounds.

HOW IT WORKS

01	02	03	04	05
Fitment quiz	Pick a tier	Sessions	Written plan	Use credits
Qualify your stack and your pain at warblecloud.com	Fixed price, from \$500	Work with Warble architects	A doc or build plan you keep	Run Starling on your clusters from day one

KEY FEATURES

- Fixed price: you know the cost before you start
- Written output on every tier
- Starling credits included: advice turns into usage on day one
- Advice from the team building Starling, Reflexion and FRAKMA
- Next step runs in your own account, with no lock-in

WHO IT'S FOR

- Founders and solo engineers choosing a first production stack
- Seed to Series B teams whose Kubernetes, GPU or MLOps setup outgrew the plan
- Mid-market and regulated teams needing custom scope (Enterprise)

PRICING

OPTION	PRICE	CREDITS	WHAT YOU GET
Niche Advisory	\$500	5,000	One deep-dive session + written recommendations doc
Growth Engagement ★	\$2,500	25,000	Three architecture sessions + a build plan you can ship
Scale Partner	\$7,500	75,000	Full architecture design + build plan, Reflexion pilot setup
Enterprise	Talk to us	Custom	Custom contract, SLA-backed, embedded with your team

★ Recommended. One-time prices. Credit top-ups: \$1 = 50 credits. Q4 Growth Partner slots: 25 Startup + 25 Enterprise seats; Startup tier price [to fill].

Status & availability. Sold at warblecloud.com. Checkout is configured in Stripe test mode; go live before promoting. Delivery times [to fill].

Not sure which tier? Take the two-minute fitment quiz.

warblecloud.com/fitment

STARLING-EX · KUBERNETES OPS VIA MCP

One command. Your cluster. Agent-ready.

Starling-EX is a self-hosted Kubernetes platform with an MCP server and CLI, so engineers and AI agents work with your cluster through scoped, auditable access, under your control.

- Stable version ready
- Integrations in development

THE PROBLEM

Simple cluster questions still need kubectl skills and five dashboards. And handing an AI agent raw cluster credentials is a security review you'd rather not have.

8

observability technologies in use per organization, on average

Grafana Labs, 2025

51%

of tech leaders report an AI skills shortage

Nash Squared, 2025

LLM01

prompt injection is the top OWASP LLM risk for 2025

OWASP, 2025

HOW IT WORKS

01

Install

Installer or Helm chart into your cluster

02

Sign in

Dex SSO, federated to your identity provider

03

Connect

Point any MCP client at Starling, or use the CLI

04

Ask

Plain-language questions, scoped tools, audited calls

KEY FEATURES

- MCP server and CLI so agents can operate Kubernetes safely
- Read-only by default for local MCP; writes need explicit opt-in **[confirm]**
- Dex SSO, Kubernetes operator and Helm chart (Community Edition, Apache-2.0)
- Runs in your cluster: credentials and workload data stay with you
- The same scoped access Reflexion uses for incident diagnosis

WHO IT'S FOR

- Platform engineers who want SSO, an operator and agent access in one install
- Teams adopting MCP-based AI agents that need safe cluster access
- Murmuration customers putting their Starling credits to work

PRICING

OPTION	PRICE	UNITS	WHAT YOU GET
Starling credits	Included	5K / 25K / 75K	Granted with Murmuration Niche / Growth / Scale Partner tiers
Credit top-ups	\$1 = 50	—	Add credits at any time
Starling-EX plans	[to fill]	—	Licence or plan pricing not yet published
Community Edition	Apache-2.0	—	Operator, Helm chart, Dex SSO

Only deck-approved prices are shown.

Status & availability. Starling-EX beta tested; stable version ready. Starling CLI v0.1.0 binaries on github.com/warbletech/starling. CE source repo is private today **[to fill]**.

Ask your cluster something in five minutes.

starling-ex.frakma.io

WARBLECLOUD AI · REFLEXION + ARMY OF AGENTS

Agents propose. Critics check. Your engineers decide.

Incident response and ops agents that run in your own account. Reflexion diagnoses incidents and drafts the fix. Shadow mode is on by default, so nothing is applied until a person approves it or your policy allows it.

Reflexion: closed beta

Army of Agents: GA after Nov 2026

THE PROBLEM

Incidents still mean someone grepping logs at 3 AM, with recovery measured in hours. Most AI ops tools that could help run in someone else's cloud, as a black box.

\$2M/hr median cost of a high-impact outage; engineers spend 33% of time firefighting New Relic, 2025	41% find outages via complaints, tickets or manual checks New Relic, 2025	97% of orgs with an AI-related breach lacked proper AI access controls IBM, 2025
--	---	--

HOW IT WORKS

01 Signal Alerts, logs, metrics, traces	02 Actor Drafts a hypothesis and a fix	03 Critic Checks evidence, policy, blast radius	04 Apply Shadow mode by default: a human approves, or policy applies
---	--	---	--

KEY FEATURES

- Reflexion: Actor/Critic diagnosis and fix proposals
- Shadow mode on by default: proposes, doesn't apply
- Every action is policy-gated, reversible and reviewable
- Agents, telemetry and models run in your VPC; cluster access via Starling MCP
- Army of Agents: specialized ops agents that observe, reason and act, on your cloud

WHO IT'S FOR

- SRE and platform teams on Kubernetes who want faster diagnosis without losing control
- Leaders who need AI ops to pass a security review: in-VPC, auditable, human-approved
- Starling users ready to move from questions to proposed fixes

PRICING

OPTION	PRICE	UNITS	WHAT YOU GET
Reflexion closed beta	[to fill]	–	Pilot setup included in Murmuration Scale Partner (\$7,500)
Army of Agents	Free to start	Pay per use	On your cloud; unlocks at GA after Nov 2026; rates [to fill]
Credit top-ups	\$1 = 50	–	Starling credits

Roadmap: Warble Brain (Avirka), Warble's own SRE models for the FRAKMA and Reflexion stack, is not trained yet.

Status & availability. Reflexion: closed beta, shadow mode on by default. Army of Agents: GA after November 2026. Portal: warblecloud.ai.

Run Reflexion on one service, in shadow mode. Measure MTTR before and after.

warblecloud.ai

FRAKMA-X · OPEN AI STACK AT SCALE

Open-source MLOps you own, with the team that built it on call.

Three layers: an Apache-2.0 Community Edition you run yourself, GitOps workflows from data to production, and an enterprise edition Warble deploys and supports on your cluster.

- Azure + GCP ready
- AWS in testing
- Apache-2.0 core
- LF wing forming

THE PROBLEM

Hand-stitched ML pipelines are hard to reproduce. Running six OSS tools together in production takes platform skills most teams lack, and managed alternatives move your data and models into a vendor's account.

5% of custom enterprise GenAI tools reach production MIT NANDA, 2025	≥50% of GenAI projects abandoned after proof of concept by end of 2025 Gartner, 2026	87% of enterprise LLM workloads run on closed-source models Menlo Ventures, 2025
---	---	---

ONE PRODUCT, THREE LAYERS

LAYER 1 · OPEN SOURCE Community Edition The FRAKMA stack: OpenFaaS, KubeRay, Argo CD, KServe, MLflow, Airflow and the WarbleApp CRD operator. Apache-2.0, always. Run it yourself. WHO IT'S FOR Data and MLOps engineers who run their own clusters; contributors. STATUS Open source. Linux Foundation wing forming; not an LF or CNCF project. Public repo [to fill]. PRICING Apache-2.0. You pay your own cloud.	LAYER 2 · WORKFLOWS GitOps MLOps Data → train → register → serve → monitor, all declared in Git. Drift triggers retraining; LLM providers swap with one manifest. WHO IT'S FOR Teams moving from notebook-driven ML to reproducible production pipelines. STATUS Documented on frakma.io; set up by Warble in the enterprise tiers. Reflexion: closed beta, shadow mode. PRICING Free to run yourself; Warble setup included in enterprise tiers.	LAYER 3 · PRODUCTION Enterprise edition FRAKMA fully managed on your AKS or GKE cluster. Warble deploys, configures and supports the platform. WHO IT'S FOR Teams running production ML and LLM workloads. STATUS Offered on frakma.io/enterprise.html; Stripe checkout for Starter and Growth. PRICING Starter \$5,000/mo · Growth \$15,000/mo · Enterprise custom.
---	--	---

HOW THE GITOPS WORKFLOW RUNS

01 Bootstrap Argo CD app-of-apps wired to your Git repo	02 Pipelines Airflow DAGs with GitSync, wired to your data	03 Train KubeRay jobs and fine-tuning pipelines	04 Register MLflow; artifacts in Azure Blob, GCS or S3	05 Serve KServe InferenceService, canary rollouts	06 Monitor Prometheus and Grafana; drift → retrain
--	---	--	---	--	---

Foundation status. Per [warbleoss.org](#), Frakma-X is Warble's Linux Foundation wing, still forming. The CNCF Sandbox path is for Starling EX Community. Neither is foundation-accepted. Which project goes to CNCF: **[confirm]**.

Run the Community Edition this week, or let Warble run it for you.

[frakma.io/enterprise.html](#)

FRAKMA-X · ENTERPRISE EDITION

FRAKMA, deployed and supported on your cluster.

FRAKMA is open source. You're buying the expertise and operational support. Every plan includes cluster setup, platform deployment and ongoing support.

Offered today

AKS · GKE

5 working days to a running cluster

PLANS · PER [FRAKMA.IO/ENTERPRISE.HTML](#)

	Starter FROM NOTEBOOKS TO PIPELINES	Growth MOST POPULAR · PRODUCTION LLM	Enterprise WHITE-GLOVE · REGULATED
PRICE	\$5,000/mo	\$15,000/mo	Custom
TERMS	Monthly · cancel any time (30 days' notice)	Monthly · 3-month minimum	Annual contract · custom terms
PLATFORM	FRAKMA deployment and configuration on your AKS or GKE cluster	Everything in Starter	Everything in Growth · dedicated AKS/GKE cluster or on-premises · multi-cluster, multi-region
AGENTS	Up to 3 WarbleApp agent deployments	Unlimited WarbleApp agents · Reflexion engine deployment and LLM provider config	Private Helm chart registry and operator builds
PIPELINES	Airflow DAG setup (up to 5) · MLflow tracking with Azure Blob artifacts · TLS, ingress, cert-manager · Prometheus and Grafana	Custom Airflow DAGs (up to 15) · KubeRay fine-tuning pipeline · KServe InferenceService and canary config	SOC 2 / HIPAA / GDPR advisory support
SUPPORT	Monthly architecture review · email, 48h response SLA	Dedicated Slack channel · bi-weekly architecture and roadmap sessions · 4h SLA (business hours)	Custom SLA incl. 24/7 on-call · platform engineer embed · quarterly exec reviews · training · custom billing
NOT INCLUDED	Custom agent development · dedicated Slack · SLA under 4h	24/7 on-call · on-premises deployment	—
HOW TO BUY	Stripe checkout	Stripe checkout	Talk to Avinash

ONBOARDING: SIGN-UP TO RUNNING CLUSTER IN 5 WORKING DAYS

01 Purchase	02 Discovery	03 Assessment	04 Deploy	05 Pipelines	06 Handoff
Kickoff link within 2 business hours	60-min call with Avinash	Written cluster review first	Full stack, TLS and DNS	Agents, DAGs, MLflow wired	Runbook, diagrams, walkthrough

INCLUDED IN EVERY PLAN

- Cluster architecture design
 - GitOps pipeline setup (Argo CD)
 - LLM agent deployment (Vertex, Groq, Ollama)
- Observability incl. LLM token costs
 - Secrets and wildcard TLS
 - Image registry and CI/CD (ACR, GitHub Actions)

OWNERSHIP AND ACCESS

- Scoped service account; everything handed back in your Git repo
- No infra markup: you pay your cloud directly
- Cancel and run it yourself: Apache-2.0 core

Before publishing: these prices aren't in the pitch deck · the Stripe product descriptions differ from the page (node limits, multi-cluster in Growth) · Starling-EX isn't in any tier **[to fill]** · Reflexion is in closed beta with shadow mode on by default; say so wherever it's sold.

Start with a 20-minute intro call or a technical deep-dive.

cal.com/frakma.io