AI on Both Sides of the Screen

Changing How We Build Products — and How We Work

Lari Suomalainen · Liana Technologies

FAST Sprint Meeting · Rovaniemi · 12 August 2026

Liana in 60 Seconds

- Finland's largest cloud service provider for **marketing and communications technology** — born in Oulu, offices in 5 different countries (CE, MEA and Nordics covered).
- **3000+ customers**, from SMEs to Tier 1 enterprises, +20k users
- ~€20M revenue · 150 people (Ilkka Group total: ~€40M · 330 people)

Not a lab. Not a startup. A mid-sized European product company with real legacy and real regulators.

Lari Suomalainen

COO, Liana Technologies

Me in 30 seconds

Over **20 years** in the Marketing Technology business

Current Role

COO at Liana Technologies

Industry

MarTech & SaaS

Communities

Member of IAB Martech Group Finland, member of SW4E ecosystem

The Frame: Two Screens

Side A — The product screen

AI our customers use. Touches customer databases. Highest governance stakes.

Side B — The developer screen

AI we use to build. Touches code. Lower governance stakes.

The mistake most companies make is treating these as one problem. They are two different risk categories and they must be governed separately.

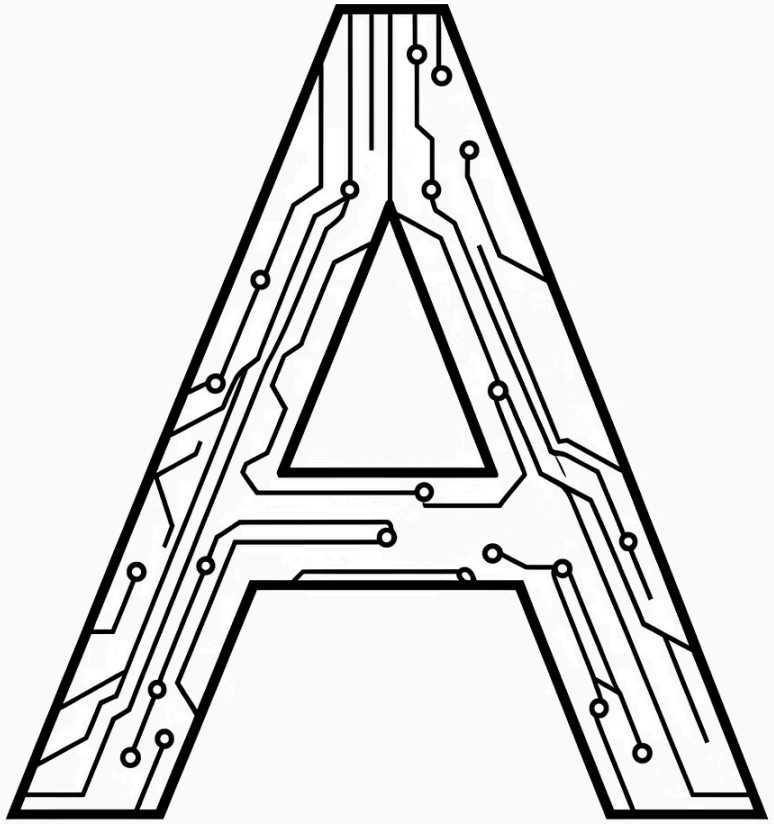

Side A

Building AI Into a Martech Platform

The Product Bet: LianaAI

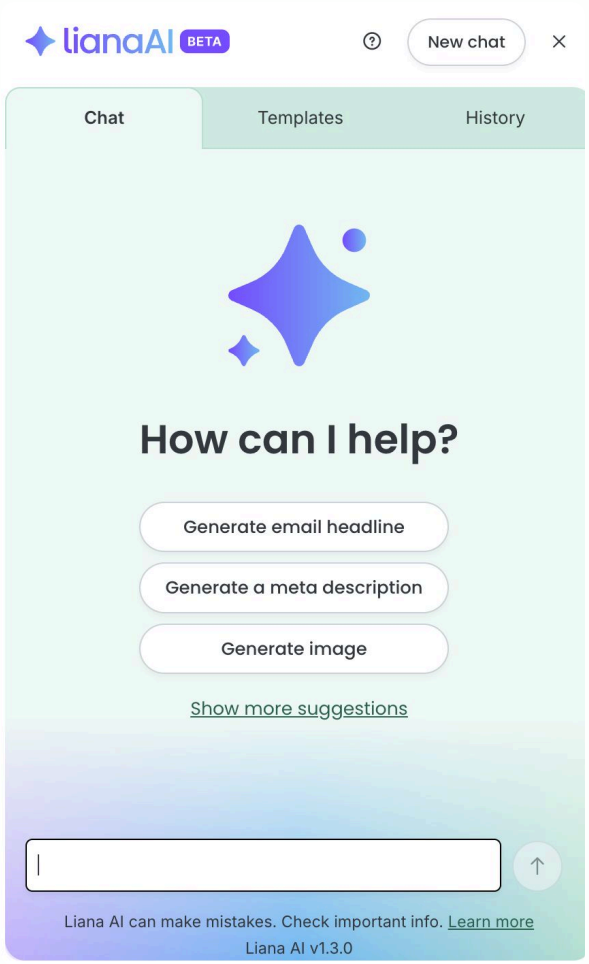

Liana Companion - "LianaAI". It´s mission is to answer to the question every marketer starts with: *What should I do?*

Design promise: **transfers "I think" into "I know."**

- Follows the user everywhere in Liana products - top panel, and at the exact spot where content is created
- Knows the customer: business landscape, own products and services, tone-of-voice
- Combines Liana product data to generate better-converting content
- **Acts only on demand.** The customer is always in control.

Why Not Just ChatGPT?

A generic assistant has the model. It does not have the context.

What LianaAI has that a general-purpose chatbot does not:

- The customer's own brand tone-of-voice
- Learning from how that customer's team actually works
- Continuity across the whole product suite, no tab-switching
- Customer-specific data protection and EU residency
- Recommendations grounded in the customer's own marketing analytics
- Marketing and communications data unified in one view

The value is not the model. It is the context pipeline around it.

Capability stages

Four phases of developing LianaAI

Phase 1 — Easy AI

1

Assistant in the top panel; text and image generation in multiple languages.

2

Phase 2 — Exact AI

Customer context added via a wizard, company site, competitors, markets etc...

Phase 3 — Insightful AI

3

Liana product data drives it automatically; analysis of past marketing campaigns and releases

4

Phase 4 — Dynamic AI

Generates straight into product UIs, for instance a full marketing email in the customer's brand and tone, adapted per channel and recipient

Quick demo

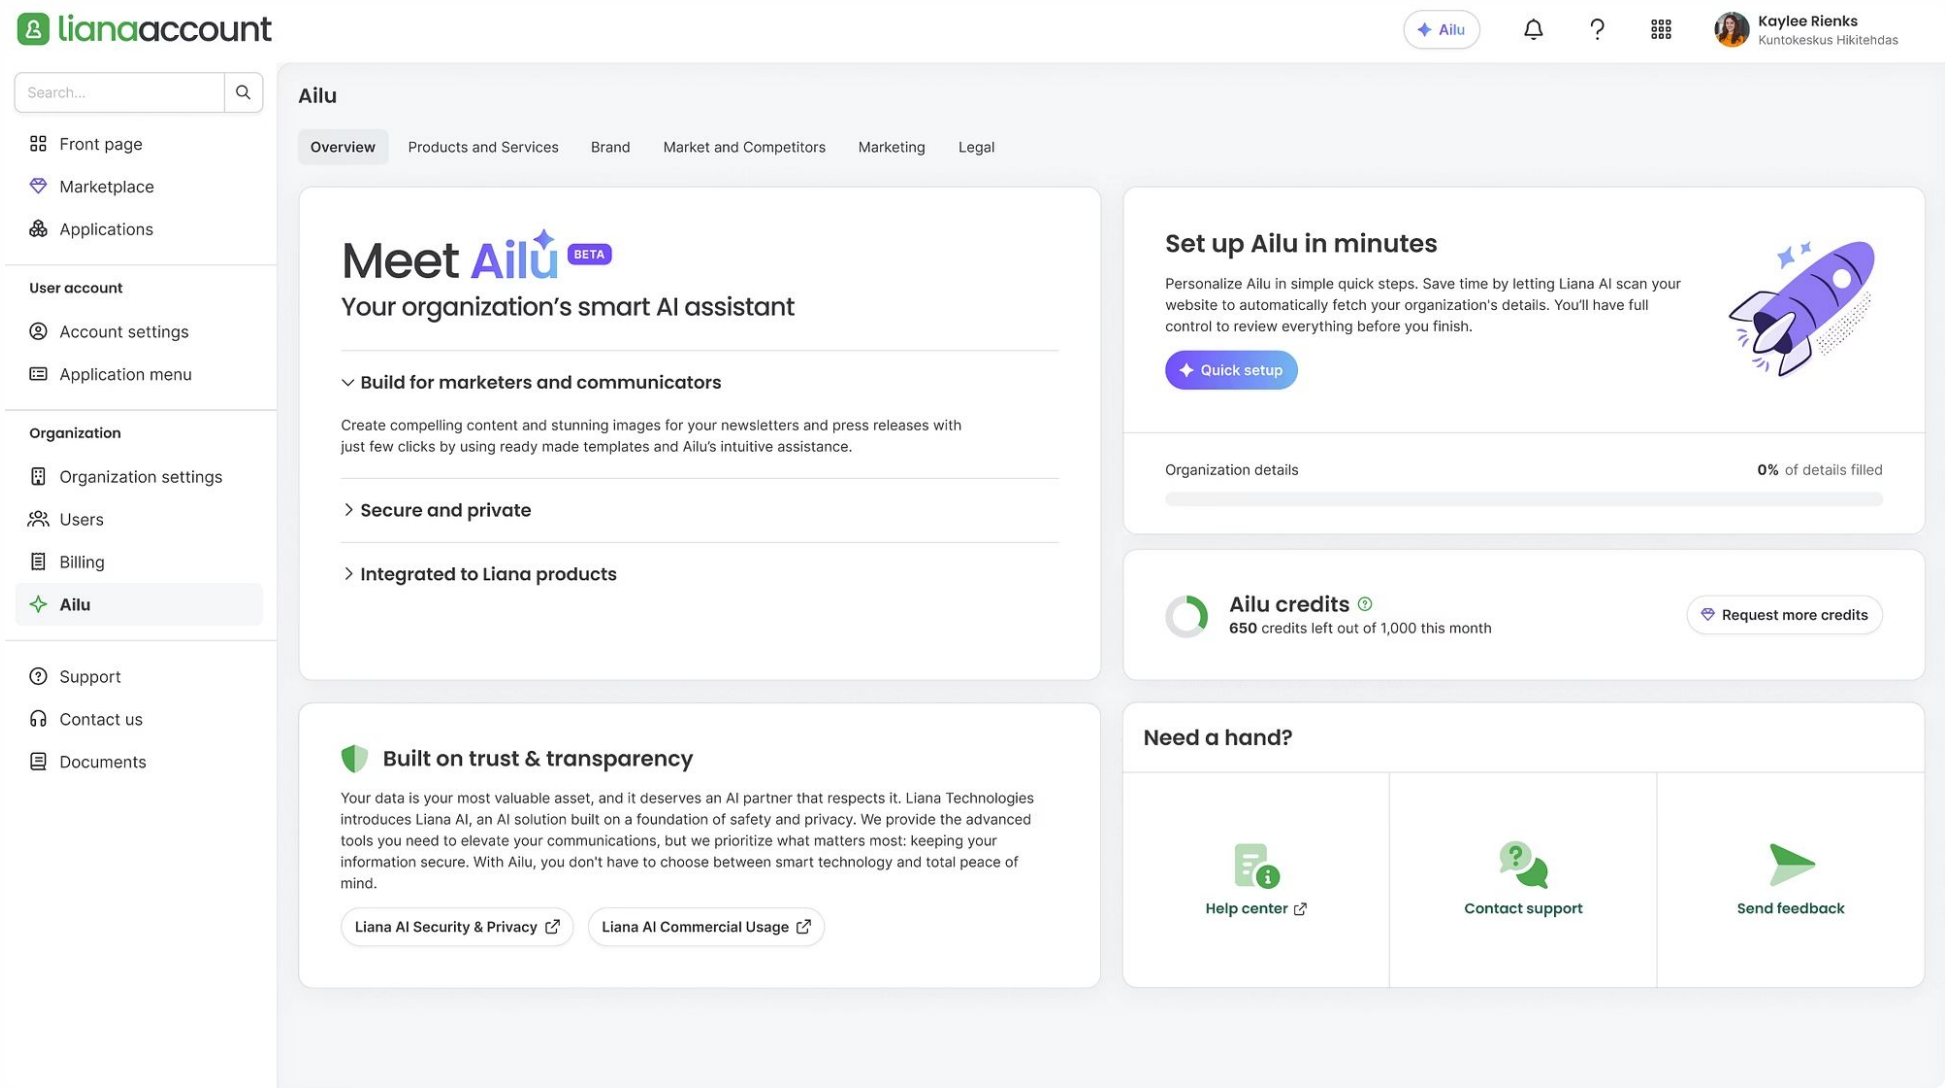
lianaaccount
Ailu
Kaylee Rienks
Kuntokeskus Hikitehdas
Search...
Front page
Marketplace
Applications
User account
Account settings
Application menu
Organization
Organization settings
Users
Billing
Ailu
Support
Contact us
Documents
Ailu
Overview
Products and Services
Brand
Market and Competitors
Marketing
Legal
Meet Ailu BETA
Your organization's smart AI assistant
Build for marketers and communicators
Create compelling content and stunning images for your newsletters and press releases with just few clicks by using ready made templates and Ailu's intuitive assistance.
Secure and private
Integrated to Liana products
Set up Ailu in minutes
Personalize Ailu in simple quick steps. Save time by letting Liana AI scan your website to automatically fetch your organization's details. You'll have full control to review everything before you finish.
Quick setup
Organization details
0% of details filled
Ailu credits
650 credits left out of 1,000 this month
Request more credits
Built on trust & transparency
Your data is your most valuable asset, and it deserves an AI partner that respects it. Liana Technologies introduces Liana AI, an AI solution built on a foundation of safety and privacy. We provide the advanced tools you need to elevate your communications, but we prioritize what matters most: keeping your information secure. With Ailu, you don't have to choose between smart technology and total peace of mind.
Liana AI Security & Privacy
Liana AI Commercial Usage
Need a hand?
Help center
Contact support
Send feedback

20 Open Initiatives, Four Pillars

AI Product Platform — 7 initiatives

Mastra agent framework + UI, file upload, token limits, Insightful AI, Dynamic AI, commercialisation of AI credits.

Data Infrastructure — 4

Data Lake POC, Lakehouse, Data Readiness Framework.

Compliance & Governance — 4

EU AI Act Article 50 content labelling, AI sovereignty, AI security, AI agent governance.

Market & Community — 5

European model test with Tampere University, structured beta programme, AI survey and whitepaper, AI community, LUT AI course POC.

Note the ratio. Only about a third of the work is "the AI." The rest is data, law, and adoption.

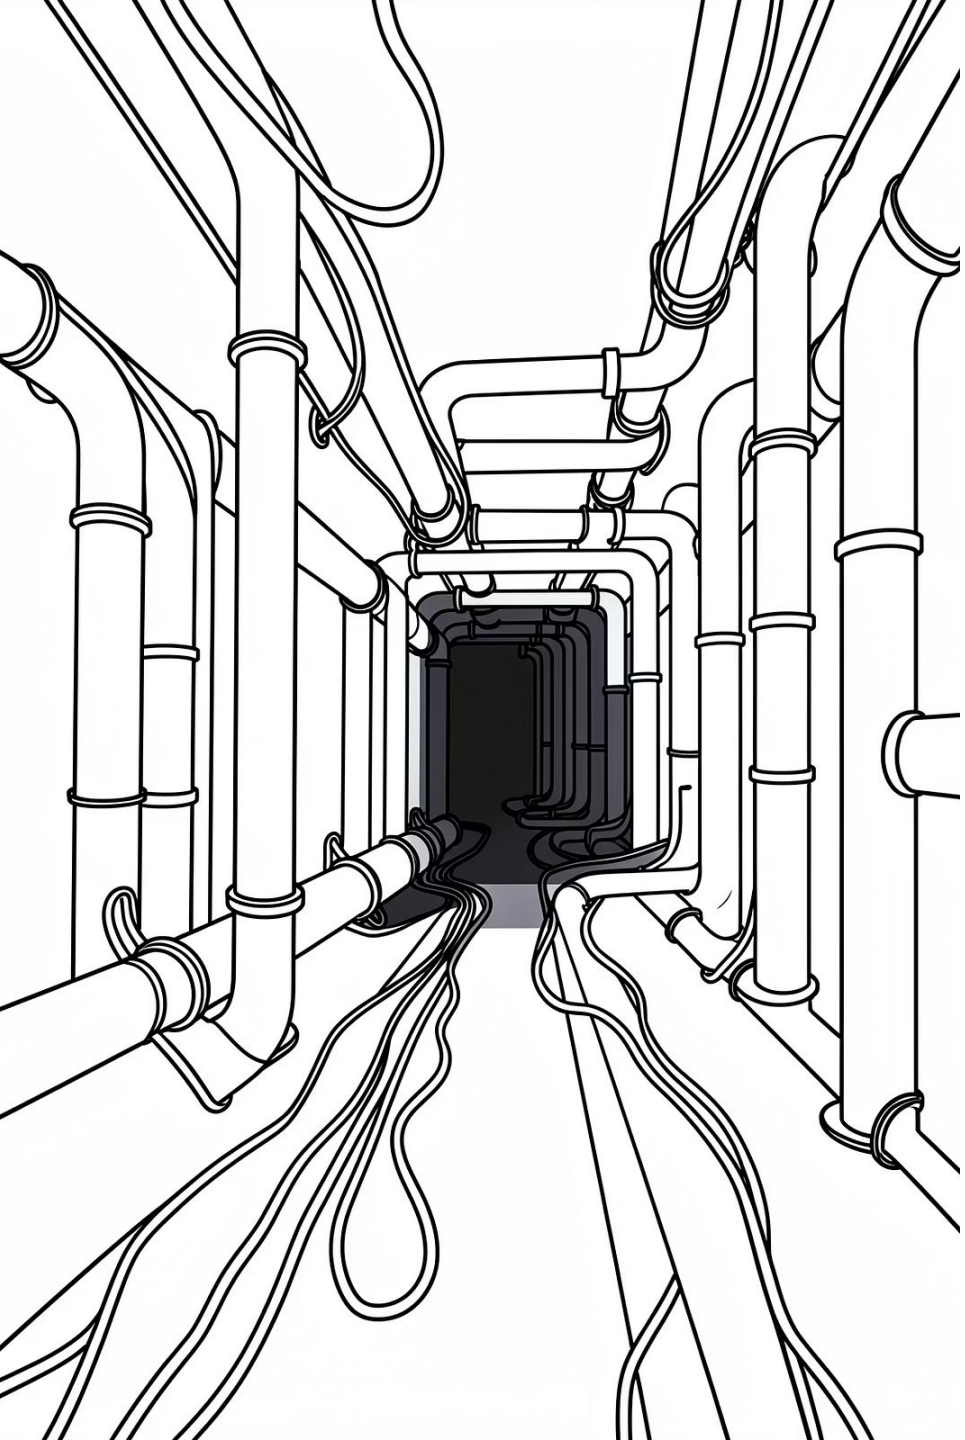

No AI Without Data Readiness

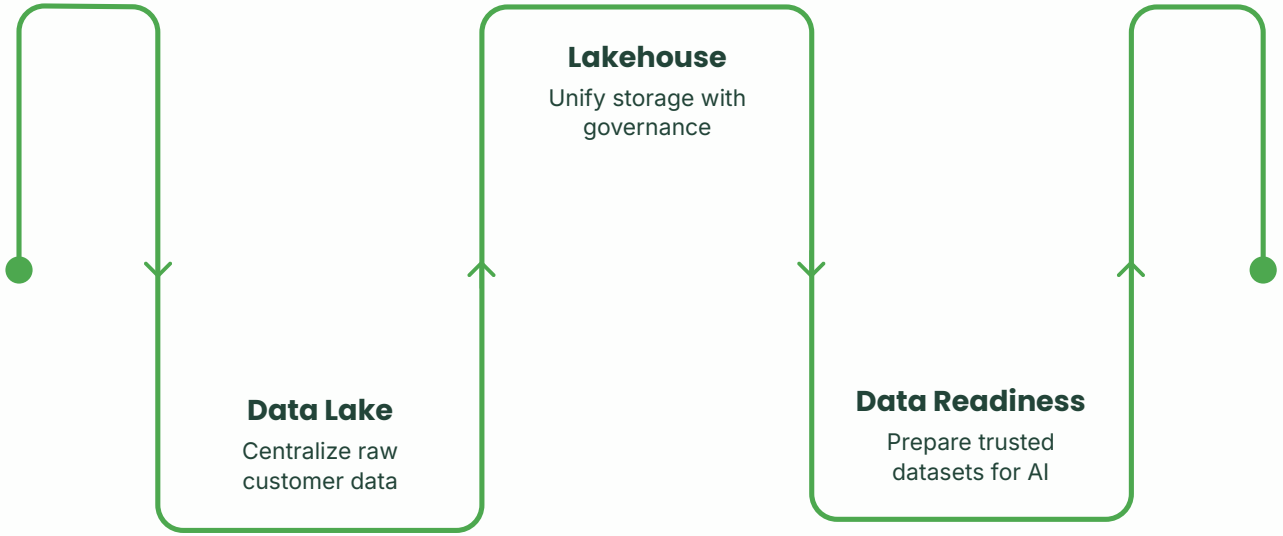

LianaCloud provides a 360° marketing view: the customer's data across Liana products *and* third-party solutions, in one place.

The unglamorous truth: the assistant is only as good as the pipeline underneath it. **Most of our AI budget is data engineering.**

Architecture: The Vertex AI Switchboard

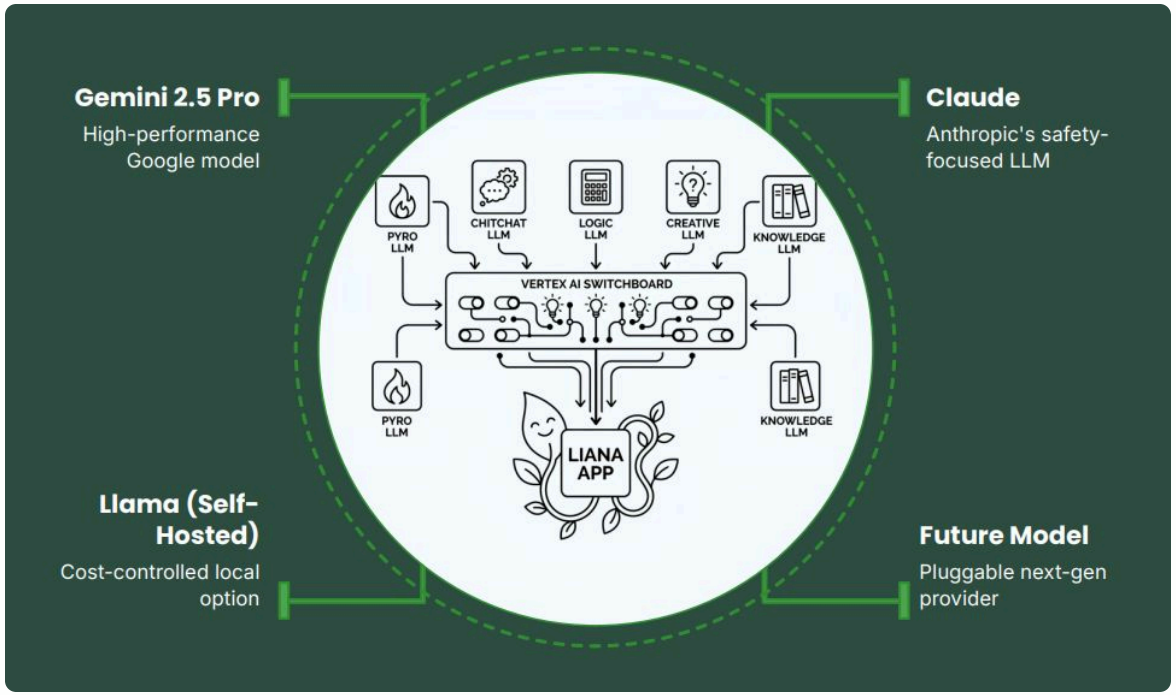

Agnostic by design. One parameter change in the API swaps Gemini for Claude or Llama. Think of it as a retro telephone switchboard: one central panel routing each task to whichever model is optimal, all surfacing through a single Liana interface.

How we got here: 2022–23, OpenAI was the only viable partner. Then Gemini 2.5 Pro levelled the field, and integrated T&Cs inside our existing Google ecosystem made it the corporate default. Claude joined after that.

Legal friction, not benchmarks, decided the default. Model choice is a runtime decision, not an architectural commitment.

Cost control

Move to self-hosted models in the Model Garden if API prices spike

Vendor leverage

No lock-in means negotiating power with every provider

Compliance as Product Architecture

EU AI Act Article 50

AI-generated content labelling and metadata, built into the product rather than bolted on afterwards.

Data Residency

Google Hamina (Finland) and Ireland. GDPR becomes a sales weapon, not a cost centre.

The Counter-Example

Our Swedish subsidiary Ungapped runs on traditional data centres. That constraint is exactly the argument for cloud-native EU infrastructure.

The Edge

We can expand AI features without legal friction from new sub-processors. Hamina means freedom to move fast.

The IP Myth: Source Code vs. Data

Our source code is **not** our most valuable asset. If it leaks, we do a 30-minute post-mortem and move on.

The real asset: customer data, brand trust, and operational logic. That is what cannot be rebuilt in an afternoon.

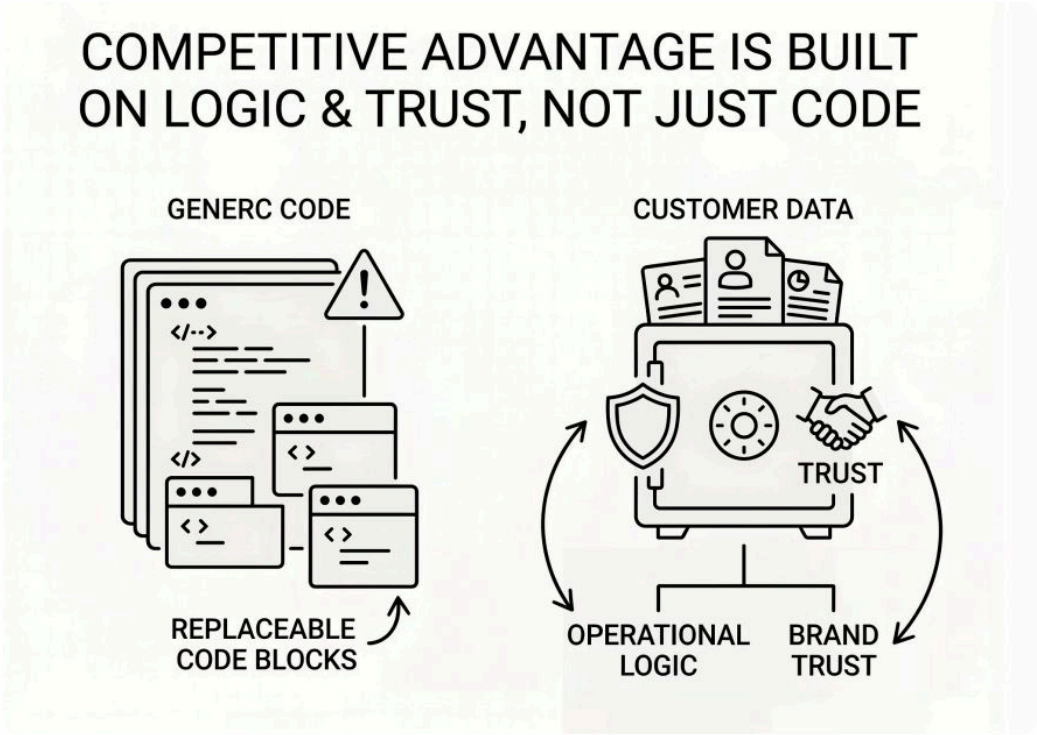

Side B

How the Work Changes

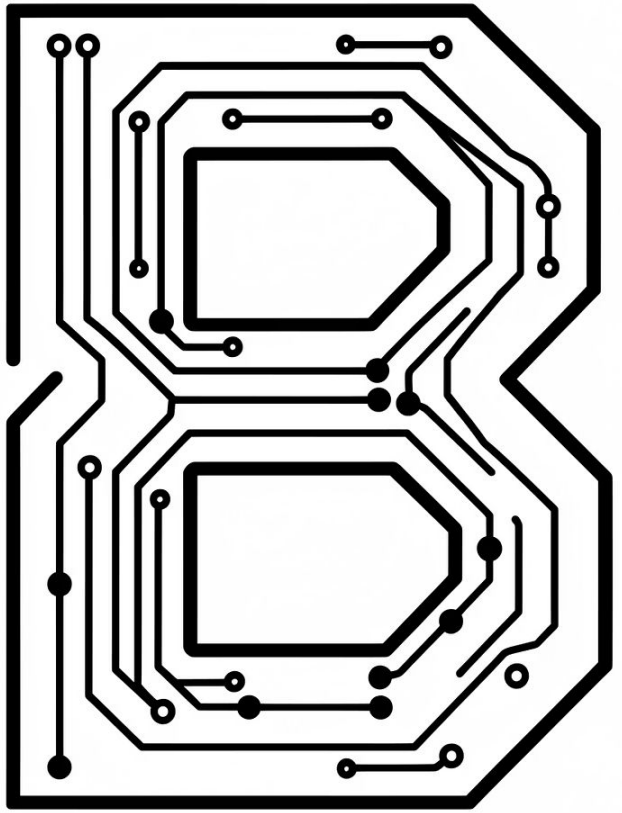

The Fail-Fast POC Culture

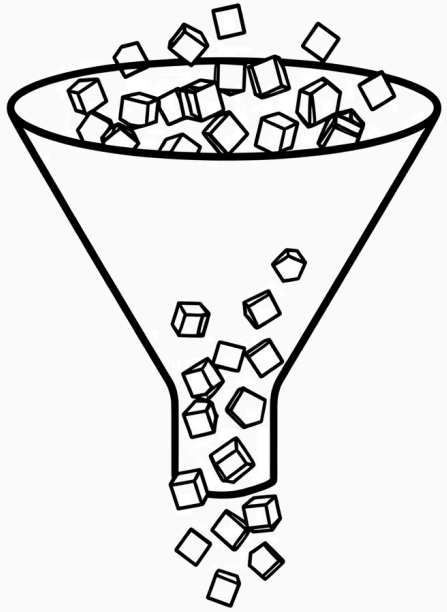

Low barrier to entry, fast barrier to exit. We POC tools weekly. Most are discarded.

100 ideas enter. A few reach full-scale adoption. Every exit is a data point, not a failure.

A "failed" POC is a win — it means we did not spend six months on the wrong technology.

For researchers: industry rarely runs clean experiments. We run cheap ones, fast, and act on noisy evidence. That gap between rigour and speed is itself worth studying.

The 2026 Developer Stack

GitHub Copilot

The reliable daily driver for boilerplate and syntax. Always-on productivity layer.

Claude + Cursor

Primary environment for deep architectural work and context-heavy codebase navigation.

GitLab Duo

Built-in repository intelligence. Honest verdict: useful, still maturing.

CodeRabbit

Automated review as a second quality gate in CI/CD. Human-in-the-loop by design.

Documentation is AI-first too. Our ISO/IEC 27001:2022 certification required 30+ pages of policy plus a Statement of Applicability with roughly 9,000 context-dependent analysis questions. All of it produced with AI.

The Trap: Crutch vs. Exoskeleton

The Crutch	The Exoskeleton
Prompt → paste → pray. You get an answer, and you do not understand the reasoning behind it.	Using AI as a hydraulic suite to lift 1,000 kg of boilerplate while ***you*** supply the brainpower. **The goal is not to avoid thinking. It is to think at a higher level.**

We don't hire coders. We hire problem solvers who can spot when the robot is lying.

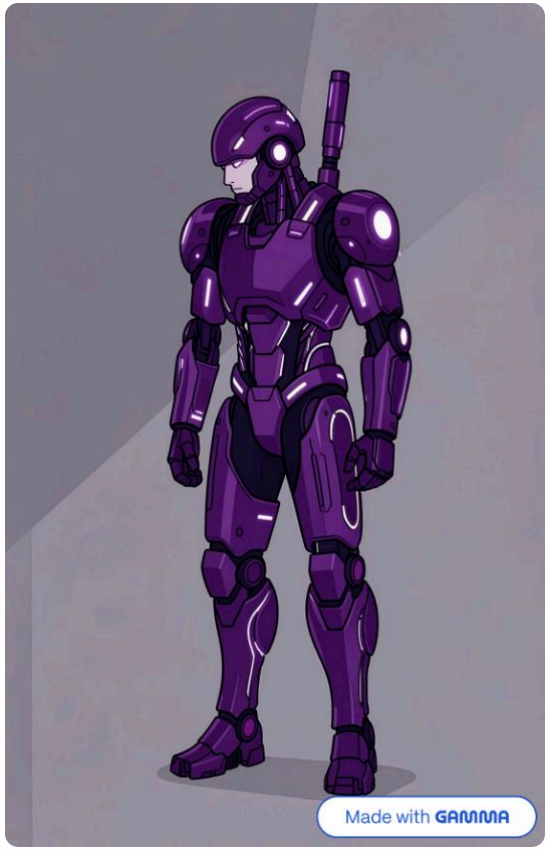

Every Line Is a Liability Until Verified

Shift in workload

The time you save on typing must be reinvested in code review and manual testing. Velocity goes up; cognitive responsibility stays exactly the same.

Confidently wrong

AI generates code that *looks* flawless while hiding subtle, costly flaws. Confidence is not correctness.

The 20/80 ratio

Spend 20% of your time generating and 80% validating. Shift the balance deliberately.

Bad tests are worse than no tests

`expect(true).toBe(true);` passes and proves nothing. AI can produce technical debt at 10× speed.

You are still the **compiler of truth**. The agent produces a draft. You decide whether it ships.

The Comprehension Debt Problem

The more you let AI write, the weaker your mental map of the codebase becomes.

The risk

Without that map, finding buggy code and reasoning about edge cases becomes dramatically slower, and more dangerous.

The scale

500k-line monoliths, ten years in production. There is no clean tour.

Our counter-moves

Semantic search by intent rather than keyword ("how do we handle authentication across these three microservices?"), and AI-generated architecture diagrams and living documentation to rebuild the map we are losing.

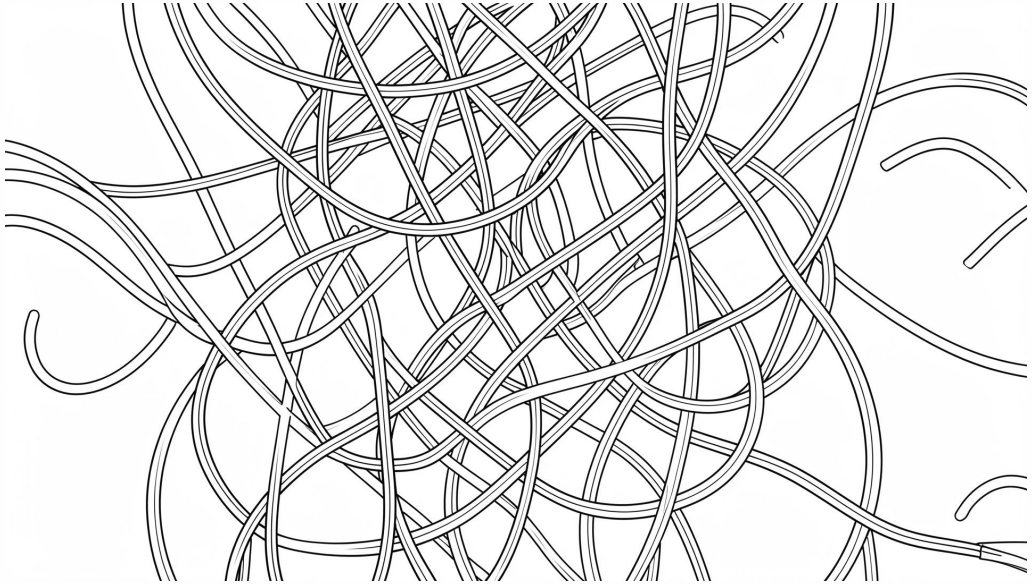

This is, I think, the most under-measured cost in the industry right now.

Architect, Not Bricklayer

1 2 3

1. Sparring Session

Use AI brainstorming to turn raw thoughts into structured Markdown specifications.

2. Technical Spec

Lock it down completely before a single line of code. No spec, no build.

3. Robot Executes

Hand over a precise, well-scoped prompt. Output quality mirrors spec quality.

Every spec defines three things: the **data contract** (exact inputs and outputs), the **side effects** (every DB write, webhook, notification — nothing implicit), and the **edge case map** (use AI to brainstorm failure modes *before* writing logic).

Syntax is cheap. Logic is expensive.

Product Intelligence: Garba

Garba is an AI-powered customer intelligence tool that automatically analyses sales conversations to turn customer feedback into actionable product insights.

Predicting Churn & Features

Identifying missing features, predicting churn, and mapping customer sentiment at scale through deep conversational analysis.

Extracting Signal

AI analyses thousands of customer conversations to surface intricate patterns and unspoken needs, often invisible to the human eye.

Sales → Dev Pipeline

Transforming qualitative "Sales Hype" into quantified "Product Specs" — precisely bridging the gap between customer input and engineering deliverables.

What Does Not Get Automated

Empathy

AI cannot feel the friction of a broken UI or recognise the delight of a well-placed feature. Only developers who use the product can advocate for the people who will.

Negotiation

"We could build this, but it will confuse the customer" is a senior move no agent will ever make on your behalf.

Judgement about what *not* to build

Knowing when to say no remains irreducibly human.

Where Your Research and Our Problems might Overlap

Six things we cannot solve on our own:

1

Measuring comprehension debt

How do you quantify the mental-map erosion agentic coding causes? We feel it. We cannot measure it.

2

Evaluating agent-generated tests

Coverage is meaningless here. What is the metric for assertion quality?

3

Operationalising the EU AI Act

Article 50 labelling and metadata inside shipping products, not compliance PDFs.

4

Governance for AI agents

What does "human-in-the-loop" mean concretely when the loop is a CI pipeline?

5

Data readiness metrics

When is an organisation's data actually AI-ready?

6

Model-agnostic architecture

Real cost, quality, and lock-in trade-offs of switchboard designs.

We Already Work With Universities

LUT Co-op

AI course proof of concept.

Tampere University

European AI model customer test with a real beta customer.

What we can offer a researcher:

- Production scale across 3000+ customers
- Real telemetry
- A fail-fast culture that lets an experiment run in weeks rather than quarters

Three Things to Take Away

01

Govern AI by what it touches, not what it is called

Product AI and developer AI are different risk categories.

02

The model is a commodity

Context, data readiness, and legal position are the differentiators.

03

Velocity is cheap. Comprehension is not.

Whatever you save on typing, spend on verification.

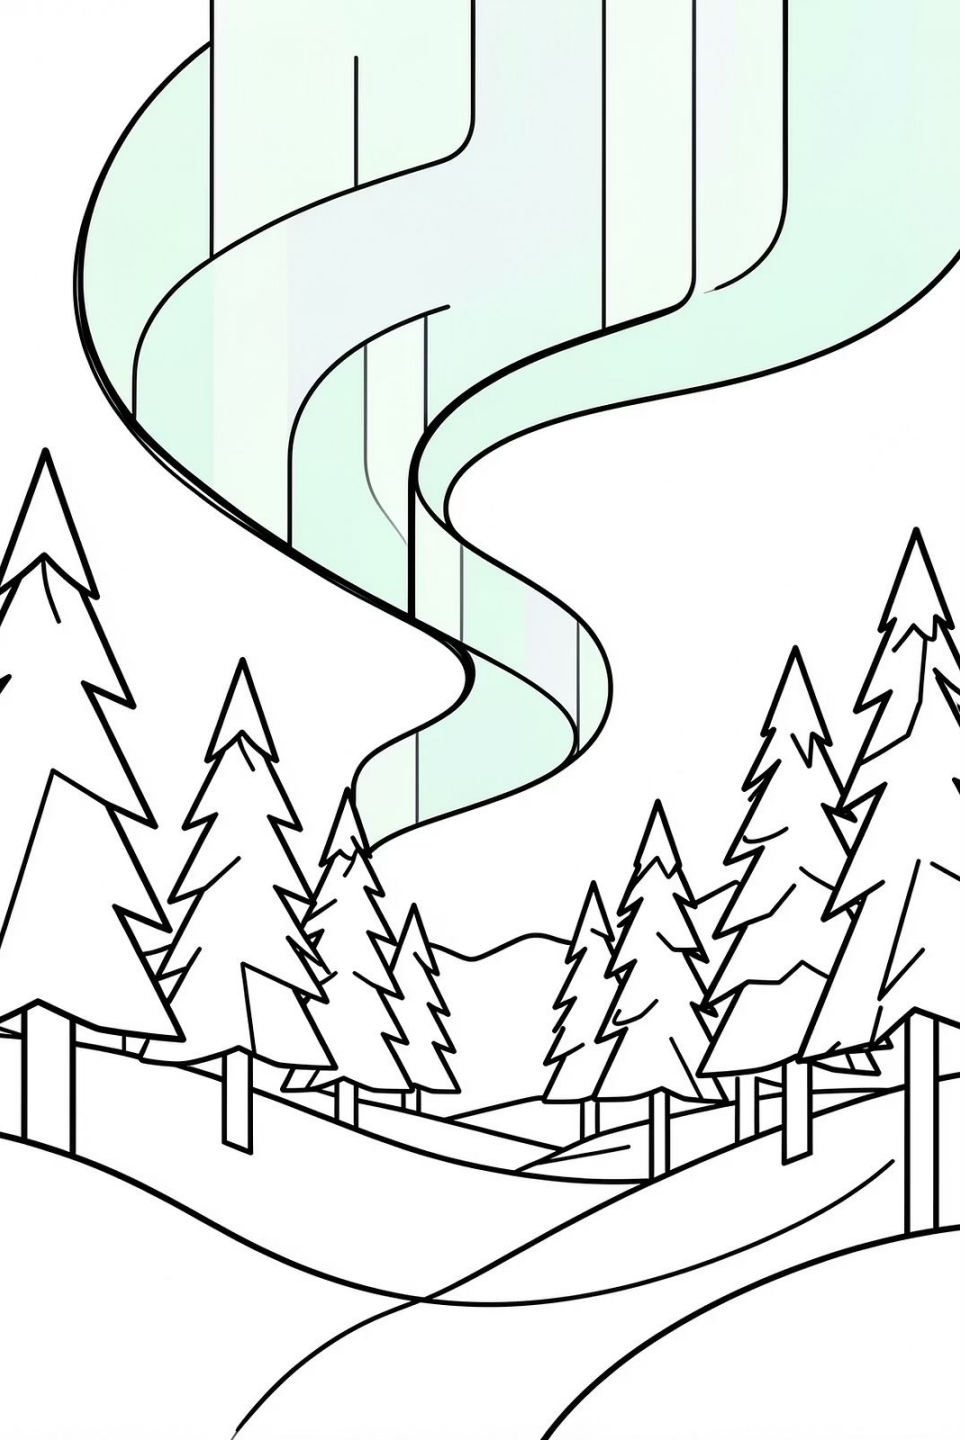

Kiitos — Questions?

Lari Suomalainen · Liana Technologies

lianatech.com/careers