

# Petru Arakiss

AI Engineering Lead. Twenty years in software, in machine learning since 2015, working on retrieval infrastructure, agent runtimes, guardrails, and evals for production AI.

Staff-level engineer with twenty years in software delivery and in machine learning since 2015. Current work centers on production AI systems that have to keep working after launch, not in a demo: document intelligence, retrieval infrastructure, agent runtimes, guardrails, eval loops, durable traces, cost and latency controls, permissions, and operator-facing product surfaces.

## TARGET SCOPE

- Principal / Staff AI Engineer
- AI Architect / LLM Platform Architect
- Forward Deployed / Solutions AI Engineer
- AI Engineering Lead

## WHAT I AM BUILDING NOW

### Current production AI systems

The names are internal; the engineering shape is the signal. Each system has a real product surface, operational constraints, and failure modes that need to be visible.

#### BIFROST

##### DOCUMENT EVIDENCE

Document intelligence for the documents finance runs on: ingestion quality gates, semantic chunking, multimodal retrieval, pgvector/HNSW search, caching, source quality, analytics, and honest no-answer behavior.

Python / FastAPI / Postgres / pgvector / Docling

#### ORVIAN

##### AGENT RUNTIME

AI workflow runtime with multi-tenant APIs, context assembly, durable memory, deterministic/cached/full-LLM tiers, run events, idempotency, queues, and human-review metadata.

TypeScript / Hono / Drizzle / Postgres / queues

#### Polaris

##### INTERNAL SEARCH

Internal AI assistant product that integrates BIFROST retrieval with MONARCH guardrails, cached safety-to-retrieval handoff, citations, streaming UX, analytics, and suggestion revalidation.

Next.js / Vercel AI SDK / BIFROST / MONARCH

## 20 yrs

in software, and in machine learning since 2015, across product, frontend, backend, and platform work

## 3 systems

current AI systems across document intelligence, agent runtimes, and internal assistants

## 12-20

engineers led in banking and fintech delivery contexts

## RECENT OWNERSHIP

### Atlax360

#### AI Engineering Lead / Head of Frontend Platform

Madrid / Sep 2021 - Present

- Lead production AI engineering across BIFROST, ORVIAN, and Polaris with high autonomy.
- Design retrieval infrastructure, agent runtime patterns, guardrails, context assembly, evals, and traceable execution paths.
- Own observability signals around latency, token and cost metadata, confidence, source quality, and human-review states.
- Translate complex backend and AI behavior into operator-facing interfaces that reduce guesswork.
- Set frontend platform standards, component architecture, review quality, and operational UI patterns.

BEFORE CURRENT AI SYSTEMS

# Delivery depth

## Sopra Steria – Lead Architect

Apr 2019 - Sep 2021

BBVA frontend scalability with Lit and Web Components; SOLUNION React rebuild; PagoFX / Santander backend leadership for a team of 12 engineers; Bankinter Angular rewrite with DDD and micro-frontends.

## knowmad mood – Backend Lead Architect, El Corte Ingles

Jun 2018 - Mar 2019

Built Node.js and Express backend infrastructure with DDD-oriented architecture in a large retail environment.

## Profile Software Services – Frontend Lead Architect, Decathlon Spain

Jun 2017 - Jun 2018

Built a modular Angular frontend from scratch and integrated it with Symfony APIs and legacy Twig surfaces.

## Relevant Traffic / Dimensional Webs / Gran Via 80

Jan 2006 - Jun 2017

Full-stack delivery across PHP, Symfony, Laravel, Rails, AngularJS, Vue, WordPress, ecommerce, and internal tools.

PUBLIC ENGINEERING SIGNALS

# Open source

### gommage

Rust policy-as-code permission harness for AI coding agents: deterministic decisions, signed audit logs, break-glass grants, and host integration.

### traceframe

Rust local-first trace recorder for AI agent workflows: append-only run evidence, hook ingestion, ledger indexing, HTML reports, and CI gates.

### nahuali

Rust self-inspecting, auditable memory engine for AI agents: surfaces evidence, provenance, and health behind each recall so callers can see which memory to trust, with an optional Ed25519-signed tamper-evident ledger.

### vestig

TypeScript structured logging library with context propagation, observability primitives, and automatic PII sanitization.

## AI CERTIFICATION

IBM Machine Learning Professional Certificate, Mar 2026.

Credential ID P9AINTGOCM4P.

NLP, LLMs, machine learning, AI, reinforcement learning, Python, deep learning.

## AI / LLMs

OpenAI, Anthropic, Vercel AI SDK, RAG, semantic and visual retrieval, vector databases, evals, guardrails, context engineering.

## BACKEND

Python, FastAPI, TypeScript, Node.js, PostgreSQL, pgvector, Redis, Supabase, Drizzle ORM, DDD, CQRS.

## FRONTEND

React, Next.js, Vue, Angular, Lit, Web Components, Tailwind, Radix, design systems.

## INFRASTRUCTURE

Docker, Kubernetes, Terraform, Vercel, Cloudflare, AWS, GCP, OpenTelemetry.

## NODE.JS CERTIFICATIONS

Linux Foundation: LFW211 Application Development; LFW212 Services Development.

## EDUCATION

Computer and Network Systems, Instituto Salvador Allende.

Initiation in Computing, Eftimie Murgu University.

## LANGUAGES

Spanish native. English professional working proficiency.