

Dstny – AI Practice Lead

As our AI Practice Lead, you will set the technical direction for everything AI at Dstny — both how AI is built into our AI Orchestration Platform and how AI is used across the company to build better software, faster. You will define the patterns, choose the foundations and raise the bar for everyone working on AI-powered features.

This is a pure individual contributor role with no direct reports. Your impact will come from the systems you design, the standards you set, and the engineers and product teams you bring along with you. You will report to the VP Architecture and partner closely with the Cloud Architect, the UI Architect, Product, Platform Engineering and SRE.

The role is split, roughly equally, across two missions:

- Architect the AI systems that ship in our product — the agents, retrieval pipelines, model-serving infrastructure, evaluation harnesses and guardrails that turn AI capabilities into reliable, multi-tenant features for millions of users.
- Build the AI practice across the company — the patterns, tooling, training and review forums that let every engineering and product team apply AI well, safely and at the right cost.

What we expect from you...

AI Architecture

- Own the architecture of AI features inside the AI Orchestration Platform: agents and orchestration, retrieval (RAG), tool use, memory, evaluation, observability, safety and cost controls.
- Define how we serve, route and observe models — including the trade-offs between hosted frontier models, open-weights models and self-hosted inference.
- Design and codify our evaluation strategy: offline benchmarks, online experiments, regression suites and human-in-the-loop review for AI features.
- Set the patterns for prompt engineering, context management, structured outputs, function calling and agent loops, with a clear opinion on when each is appropriate.
- Write RFCs and architecture decision records (ADRs), lead design reviews, and align engineering, product and operations around pragmatic AI architecture choices.
- Partner with Security and Compliance on AI risk: data handling, privacy, prompt injection, model abuse, GDPR and customer-data boundaries.

AI Practice

- Build and lead Dstny's AI practice: the playbook of patterns, anti-patterns, reusable components and reference implementations that other teams can adopt confidently.
- Drive responsible adoption of AI in our internal engineering workflow — coding assistants, code review, test generation, documentation, knowledge search — measuring real productivity impact, not vibes.
- Run an internal AI guild / community of practice: regular forums for engineers and product to share what is working, what is not, and what we should standardise.

- Coach senior engineers, tech leads and product managers to integrate AI capabilities into their roadmaps with realistic expectations and clear success metrics.
- Track the AI landscape — models, frameworks, agentic patterns, evaluation tooling, regulation — and translate it into concrete recommendations for Dstny.
- Represent Dstny's AI practice in technical discussions with customers, partners and the wider engineering community when needed.

What We're Looking For

Education & Experience

- Master's or Bachelor's degree in Computer Science, Engineering, Applied Mathematics or a related field (or equivalent practical experience).
- 10+ years of professional software engineering experience, with several of them spent designing and shipping AI / ML systems in production.
- Demonstrated experience as a Staff, Principal or Architect-level IC — you have owned the design of AI systems and led engineers across teams without direct authority.

Core Skills

- Hands-on experience building production systems with modern LLMs: RAG, agents, tool use, function calling, structured outputs, evaluation and observability.
- Strong opinion on AI evaluation: you have built or operated evaluation harnesses (offline + online) and you know what good looks like.
- Solid grounding in distributed systems: APIs, messaging, latency, throughput, caching, backpressure, multi-tenancy.
- Production experience with Kubernetes-based platforms — including how to fit GPU / inference workloads into a cloud-native runtime.
- A pragmatic view on cost, latency and quality trade-offs across hosted and self-hosted models.
- Clear understanding of AI risk: privacy, security, prompt injection, data leakage, fairness, compliance (GDPR and similar).
- Excellent written and verbal communication: you can write a clear RFC, run a productive design review, and explain trade-offs to engineers, product and executives.
- A track record of growing engineers and shifting how an organisation works — your influence is documented, not just remembered.
- Fluent in English; Portuguese, French or Dutch is a plus.

Technologies We Use

- Backend: Python, Node.js, TypeScript, JavaScript, C, SQL, Go.
- Web apps: React, HTML, CSS.
- AI / ML: hosted LLM APIs, open-weights models, vector stores, agent and RAG frameworks, evaluation tooling.
- DevOps & QA: GitLab CI/CD, Ansible, PyTest, Jest, Robot Framework, Aikido.
- Virtualisation & containers: Kubernetes, Docker, LXC.
- Networking & security: IP networking, firewalling, Nginx, Kong, Keycloak.
- Observability: Prometheus, Grafana, Loki, Tempo.

We don't expect deep expertise in every technology on this list. We do expect that you can quickly get to a working understanding of any of them, form an opinion on the right tool for the job, and help us make sound long-

term choices. At this level, pragmatism, judgement and follow-through matter more than any single tool.

Bonus Points

- Experience building agentic systems or AI-native products at scale.
- Experience operating GPU / inference infrastructure (hosted or self-managed).
- Experience with telecom, UCaaS, CPaaS or real-time communication platforms.
- Contributions to open-source AI projects, technical writing or public speaking.
- Background in security-sensitive or regulated environments (GDPR, ISO 27001 or similar).

What we offer you...

- A senior, high-leverage role with real scope — your decisions will shape both our product and how the company builds software.
- A compensation package benchmarked for architect-level ICs, including base, bonus and benefits.
- Flexible working hours.
- Experienced colleagues, respectful engineering culture, and customers who trust us with mission-critical services.

<https://dstny.com/>