Ben Morris

linkedin.com/in/benmorr

Engineering leader who manages teams and ships production systems — built and led the engineering team behind Tenera's AI-native SaaS platform, owning everything from a multi-agent LLM pipeline to hiring and people development. Earlier, a distributed-systems engineer at Amazon Alexa, working on the orchestration layer behind every Alexa request. Looking to lead an established engineering team at a company in its growth stage.

EXPERIENCE

Tenera Inc · Toronto, ON — Sep 2021 – Present

AI-native SaaS platform for engineering & construction firms

Chief Technology Officer — *Jan 2026 – Present*

- Lead engineering and technical strategy for an AI-native SaaS platform serving ~10 paying customers, with ~50 firms in active evaluation.
- Own architecture of the end-to-end LLM pipeline powering automated report generation.
- Manage the engineering team end-to-end — 1:1s, performance, career development, and hiring decisions.
- Recovered ~$1M in non-dilutive funding (SR&ED) over four years by owning eligibility strategy, time-tracking discipline, and consultant management.

Tech Lead & Senior Software Development Engineer — *Sep 2024 – Jan 2026*

- Designed and shipped Tenera's core report-generation pipeline — LLM-driven document synthesis with end-to-end observability — cutting first-draft turnaround from a self-reported ~3-hour average to ~15 minutes.
- Built the multi-tenant data and infrastructure foundation supporting the platform at 99.99% uptime with zero cross-tenant data incidents.
- Led and grew the engineering team from 3 to 6 — including a remote hire and an external contracting partner — driving technical direction, 1:1s, and hiring.
- Drove adoption of AI development tooling across the team (Copilot, Cursor, Claude Code, CodeRabbit), substantially increasing per-developer throughput.

Founding Engineer — *Sep 2021 – Sep 2024*

- First engineering hire; built the initial product from prototype to production across frontend, backend, and AI integration.
- Shipped the early product that landed Tenera's first paying customers.

Amazon (Alexa) · Seattle, WA / Toronto, ON — Jul 2018 – Sep 2021

Second engineer on the Toronto team that grew to own Alexa orchestration end-to-end as the service was decoupled.

Software Development Engineer (L4 → L5), Alexa Orchestrator - Toronto

- Helped decompose the Alexa orchestration monolith into discrete microservices, owning the separation of core routing from skill-execution logic — work spanning every Alexa request.
- Redesigned how request context (geolocation, identity, preferences) flowed between newly separated services, replacing an implicit shared data store with explicit, entitlement-based contracts.
- Drove alignment across upstream and downstream teams to retire illegitimate data dependencies rather than preserve them, reducing superfluous processing across the majority of Alexa requests.

Software Development Engineer (L4), Alexa AI — Seattle

- Built a configurable runtime for hypothesis-ranking. Groovy compute graph evaluated at runtime, moving model swaps and input transformations from compiled code changes to plain config, enabling hot-swappable models.

• Launched Alexa's hypothesis-ranking into new international locales (India, Australia), standing up serving infrastructure and per-locale model deployment for previously unsupported markets.

SKILLS

Languages	TypeScript, Java, Python, SQL
AI / ML	LLM pipelines, RAG · Anthropic, OpenAI, Gemini · Bedrock · LangChain, LangFuse · pgvector
Cloud / Infra	AWS (Lambda, ECS, RDS, Route 53) · Terraform, Docker · GitHub Actions (CI/CD)
Data / Observability	PostgreSQL, DynamoDB · CloudWatch, PostHog, PagerDuty

EDUCATION

Bachelor of Software Engineering, University of Waterloo — Waterloo, ON · Graduated 2018