

Danish Khan

Engineering Leader | AI-Native Platforms & Distributed Systems

danishkhan.id@gmail.com | +1 (917) 562-9244 | linkedin.com/in/danishkhan

Summary

Engineering leader with 15+ years of experience building and scaling AI-native products, LLM-powered workflows, and distributed platforms in high-growth environments. Currently leads Agent Platform engineering at Highspot, defining architecture for tool execution, human-in-the-loop controls, memory, and governance across enterprise AI agents. Previously delivered GenAI-assisted content enrichment in production at Loblaw Digital, ML-powered recommendation platforms at Morgan Stanley serving 16,000+ advisors, and mission-critical financial infrastructure at Amazon. Combines hands-on engineering depth with a track record of leading 7–30+ person organizations, turning ambiguous opportunities into reliable, measurable products.

Core Strengths

AI-Native Development: LLM-powered workflows, agent architecture, RAG/vector search, human-in-the-loop quality controls, prompt engineering, production evaluation

Content & Document Generation: GenAI-assisted enrichment pipelines, rule-fused LLM workflows, structured content production at scale

Engineering Leadership: Team building and scaling, hiring, coaching, performance management, technical strategy, roadmap execution, cross-functional influence

Platform Building: Multi-tenant SaaS, workflow engines, APIs, ingestion pipelines, distributed systems, reliability, observability, cloud infrastructure (GCP, AWS)

Work Experience

Highspot, *Vancouver, Canada*

Senior Engineering Manager, Ecosystem & Agent Platform | *Jul 2025 – Present*

Lead a growing engineering team responsible for shared agent and ecosystem capabilities such as tool execution, context/memory, observability, governance, and policy enforcement, enabling three major enterprise agent families with diverse SDKs (Pydantic, AI SDK, Mastra) against a common agent infrastructure.

- Platform adopted by all three agent families within six months, eliminating redundant agent infrastructure builds and standardizing tool execution, memory, and governance across the org.
- Grew shared tool and skill registry to support 50+ tools across enterprise use cases with policy-gated access controls and human-in-the-loop approval flows.
- Established 10+ architecture RFCs and technical standards adopted as the shared baseline across all agent initiatives, reducing integration incidents and improving delivery consistency.
- Defined the data model and reporting architecture for customer-facing AI usage analytics, separating adoption insights from operational telemetry and tracking token consumption, action audit trails, and per-customer limits.
- Partner with Product, Design, GTM, Analytics, Security, and agent teams to translate evolving product needs into incremental roadmaps, clear ownership boundaries, and production-ready platform investments.

Loblaw Digital, *Toronto, Canada*

Senior Software Development Manager / Software Development Manager | *Apr 2022 – Mar 2025*

Led engineering teams and managers across search, catalog, merchandising, and data platforms serving 4M+ monthly users across three business verticals (Grocery, Pharmacy & Beauty, Apparel).

- Directed GenAI/LLM-assisted enrichment and semantic matching initiatives, combining model capabilities with deterministic workflows, quality controls, domain rules, and measurable evaluation criteria, directly analogous to AI-driven document generation workflows.
- Built and scaled Helios product-discovery platform spanning keyword and semantic retrieval, ML re-ranking, personalization, typeahead, and experimentation; reduced search failure rates by roughly 50% and improved engagement and revenue outcomes.
- Led COMET, a greenfield multi-tenant ingestion, enrichment, and merchandising workflow platform that replaced legacy SAP Hybris processes; delivered \$1M+ in annual savings, reduced merchandising turnaround by roughly 30%, and lowered enrichment defects by 25–30%.
- Partnered with Product, Design, Merchandising, Operations, and Data Science to define strategy, success metrics, experiments, and roadmap trade-offs across ambiguous, high-impact product areas.
- Owned hiring, leadership development, performance management, architecture reviews, operational readiness, and reliability practices across an organization of 30+ direct and indirect reports.

Amazon, *Vancouver, Canada*

Software Development Manager, Amazon Tax Engine | *Mar 2020 – Apr 2022*

Led 12+ engineers across tax-rule management and automated-testing platforms for business-critical financial systems, providing architecture direction, delivery leadership, and operational ownership.

- Architected TACOMA, a modular rule-management platform that replaced legacy workflows and created a scalable foundation for tax-content operations.
- Directed Hopper, a SaaS shadow-testing platform that replayed production traffic and stubbed non-deterministic dependencies, cutting testing cycles by 50%, improving coverage by 30%, and executing millions of test cases in under five minutes.
- Raised engineering quality bar through stronger design reviews, CI/CD, testing strategy, incident readiness, and hands-on technical mentorship.

Morgan Stanley, *New York, US / Mumbai, India*

Vice President / Engineering Manager / Senior Software Engineer | *Apr 2011 – Feb 2020*

Led distributed engineering teams building ML-driven recommendation and advisor-productivity platforms for Global Wealth Management.

- Designed and delivered Next Best Action, a recommendation platform for 16,000+ financial advisors that increased client engagement by 20% and contributed to \$200M+ in revenue growth.
- Built scalable ingestion, retrieval, ranking, and rule-processing systems using Java, Spring Boot, Solr, Hadoop, MapReduce, and enterprise data platforms.
- Managed 8+ engineers across US and India, delivering 5+ major releases annually with 98% on-time delivery while partnering closely with product and business stakeholders.

Education

Bachelor of Engineering (Computers) | **Mumbai University, India** | May 2005

Certification

Google Cloud Certified Generative AI Leader

Skills & Core Competencies

AI / ML: LLM applications and agent frameworks | RAG and vector search | Human-in-the-loop workflows | Context engineering | Production evaluation | Semantic retrieval | Recommendation systems | Personalization | Experimentation | PostgreSQL/pgvector | Elasticsearch

Engineering: Python | Java | TypeScript | React | Elixir | Scala | GCP | AWS | Kubernetes | Docker | Pub/Sub | BigQuery | Terraform | CI/CD } Observability | Distributed systems