

Frank Goortani

Solution Architect | AI/ML Engineer | Full-Stack Developer

Location: Aurora, Ontario, Canada

Email: frank.goortani@gmail.com

Phone: (647) 983-5277

LinkedIn: [linkedin.com/in/frankgoortani](https://www.linkedin.com/in/frankgoortani) (13K+ followers)

GitHub: github.com/FrankGoortani

Medium: [@FrankGoortani](https://medium.com/@FrankGoortani)

Website: goortani.com

Portfolio: visionzlab.com | architect.solutions

Professional Summary

Senior AI/ML Engineer and Solution Architect with **25+ years** of full-stack development experience, specializing in **production LLM systems, distributed microservices, and enterprise-scale web applications**. Currently **Full-Time CTO at FasterOutcomes**, scaling an AI legal tech startup from early revenue to Series A. Previously Solution Architect at Uber, where I built the **award-winning ELLE AI Decision Engine** (won multiple internal innovation awards) - an LLM-powered security and privacy automation platform processing millions of decisions across Uber's global operations.

Deep expertise in **LangChain, LangGraph, RAG systems, and AI Agents** with proven track record shipping production AI/ML features at scale. At Uber, led development of **9+ enterprise privacy web applications** using React 18, TypeScript, and Fusion.js, and contributed to **50+ security infrastructure projects** in Python and Go. Architected authorization middleware achieving **96% test coverage** while solving complex distributed systems challenges.

At FasterOutcomes, leading product strategy and technical architecture for production LLM applications, scaling engineering team from 3 to 15+ engineers. Active thought leader with 13K+ LinkedIn followers and technical blog on Medium. Proven ability to work across the full stack (backend, frontend, AI/ML, cloud, security) while maintaining exceptional code quality and system reliability.

Core Specialties: Production LLM Systems • AI Agents • RAG • LangChain/LangGraph • Go Microservices • React/TypeScript • Python (FastAPI/LangFx) • Distributed Systems • Cloud Architecture (GCP/AWS) • Security Engineering

Core Technical Skills

AI/ML & LLMs

- LangChain, LangGraph
- FastMCP, MCP Development
- AI Agents, RAG Systems
- OpenAI, Anthropic Claude
- Google Gemini, Cohere
- PyTorch, TensorFlow
- Prompt Engineering

Agentic Coding & AI Tools

- Claude Code, Cursor
- GitHub Copilot, Codex
- Aider, AI-Assisted Dev
- MCP Server Development
- Agentic Workflows
- 3-5x Productivity Gains

Backend Development

- Go (YARPC, gRPC)

- Python (FastAPI, Pydantic)
- Java (Spring Boot)
- Node.js, TypeScript
- Microservices, REST APIs
- Cadence, Temporal
- Apache Kafka

Frontend Development

- React 18, Hooks
- TypeScript 5.7+
- Fusion.js, NextJS
- TanStack React Query
- BaseUI, shadcn/ui
- Styletron, TailwindCSS
- Jest, Playwright

Cloud & DevOps

- Google Cloud (GCP)
- AWS, Azure
- Docker, Kubernetes
- Bazel, Helm
- GitHub Actions, Jenkins
- Terraform, IaC
- M3, Prometheus

Databases & Data

- Cassandra, MySQL
- PostgreSQL, MongoDB
- Redis, DynamoDB
- OpenSearch, Pinecone
- Qdrant, Chroma
- Apache Spark, Presto
- Kafka Streams

Security & Identity

- Charter (Uber IAM)
- SPIFFE/SPIRE
- OAuth2/OIDC, JWT
- RBAC, authfx
- GDPR/CCPA
- KMS, PKI
- Security Engineering

Professional Experience

Uber Technologies | Solution Architect (Contract via BayRockLabs)

January 2021 - December 2025 (5 years) | Remote

Leading AI/ML and full-stack development initiatives across Uber's Technical Privacy and Engineering Security organizations. Built award-winning production LLM systems, architected enterprise web applications, and contributed to 50+ security infrastructure projects.

ELLE (EngsecLLM Engine) - AI Decision Engine - Primary Project

Award-winning LLM-powered security and privacy automation platform

Achievement: Won multiple internal Uber innovation awards for automating Technical Privacy review workflows, reducing manual review time by 70% while improving decision consistency.

Technical Leadership:

- **Architected and built** production LLM system processing millions of security/privacy decisions across Uber's global platform
- **Designed multi-agent architecture** using LangChain and LangGraph for self-correcting AI workflows:
 - Processor Agent + Reviewer Agent pattern with iterative improvement

- RAG (Retrieval Augmented Generation) with OpenSearch vector store integration
- Template-based prompt management for consistent decision-making
- Structured output parsing with Pydantic for reliable data extraction
- **Implemented 6+ automated workflows:** Technical Privacy Review (L1/L2 data classification gates), Security Guidance automation, Legal Review routing, Privacy Impact Assessment, Community Operations Review, JIRA prioritization (automated P0/P1/P2 assignment based on risk analysis)
- **Built Google Docs ingestion pipeline** for document RAG with semantic search and metadata filtering
- **FastAPI service architecture** with /agent and /agent-for-template endpoints
- **Integrated with Uber platform:** M3 metrics, Flplr feature flags, JIRA automation, Salesforce
- **Designed authorization middleware** (Go, YARPC, Charter) with 96% test coverage:
 - Dual middleware architecture (inbound/outbound) for comprehensive endpoint coverage
 - Eliminated circular dependencies using MCP Gateway pattern with fx.Lifecycle hooks
 - LDAP group-based access control with Charter Decision Service integration
 - SPIFFE ID extraction and employee ID mapping

Technologies: Python (FastAPI, LangChain, LangGraph, FastMCP, Pydantic), Go (YARPC, gRPC, Protocol Buffers), OpenAI API, Anthropic Claude, OpenSearch (vector store), Cadence workflows, Cassandra, Redis, M3 metrics, Bazel, Claude Code (agentic development)

Business Impact:

- Automated 70% of privacy review decisions, reducing review time from days to minutes
- Enabled 6+ security use cases with consistent LLM-powered decision-making
- Centralized GenAI platform serving multiple Engineering Security teams
- Charter audit trail for compliance and security governance
- Achieved 3-5x development velocity leveraging Claude Code and agentic coding workflows

Enterprise Privacy Web Applications (9 Production Applications)

Full-stack development using React 18, TypeScript, Fusion.js

Led development of 9+ enterprise privacy web applications for Uber's Technical Privacy team, serving thousands of internal users with production-grade security, comprehensive testing, and Uber platform integration.

Applications:

1. **Elle Web** - EngsecLLM Engine Frontend (LLM-powered privacy review automation UI, ERD submission workflows, real-time decision flow, JIRA integration)
2. **IRIS Web** - Privacy Web Frontend (Privacy request management, multi-region GDPR/CCPA compliance, GeolIP integration, Presidio WebView)
3. **DSAR Web** - Subject Access Request Automation (GraphQL-based data fetching, schema stitching, Privacy Guard Challenger integration)
4. **Consents Manager Web** - Contextual consent management (Consent resource creation, management key generation, RPC-based microservice architecture)
5. **Web CMP** - Consent Management Platform (IAB-approved TCF v2 implementation, cross-border consent management, Rosetta internationalization)
6. **TPRM QC Copilot Web** - Third-Party Risk Management (LLM-powered quality control for vendor risk assessment, CSV data parsing)
7. **Agentic AI Observability Web** - LLM observability and monitoring platform
8. **Design Review Web** - Privacy design review workflow automation
9. **Privacy Playground** - Privacy engineering experimentation platform

Cross-Application Technical Excellence:

- **Comprehensive testing:** Jest (34+ test suites), Playwright (E2E), React Testing Library
- **Type-safe RPC:** Protocol Buffers with automated TypeScript code generation
- **Universal rendering:** Fusion.js SSR with client hydration, code splitting, performance optimization

- **Platform integration:** 30+ Uber services (USSO auth, Flipr feature flags, M3 metrics)
- **Security:** CSRF protection, CSP, JWT sessions, secure headers, production-grade security
- **Observability:** M3 metrics, structured logging, distributed tracing, APM
- **Build system:** Bazel monorepo, Yarn Workspaces, TypeScript project references
- **Design system:** BaseUI compliance, Styletron CSS-in-JS, responsive design, accessibility (ARIA, keyboard navigation)

Security Infrastructure Projects (50+ Python Projects)

Enterprise-scale Python monorepo across AI/ML, data engineering, and security

AI-Powered Security (AI CoE - Center of Excellence):

- **SecureCode AI:** Privacy-aware static code analysis for PII detection in logs (LLM-based code scanning, VSCode extension, regex + LLM hybrid detection)
- **Agent Guardrails:** Security controls for AI agents (permission checking, reflection-based guardrails, MCP integration)
- **Custom MCP Servers:** Built Model Context Protocol servers for agentic workflows and tool integrations (FastMCP, Python)
- **Container Image Remediation Assistant:** LangFx Service with LangGraph (AI-powered vulnerability remediation, JIRA integration)
- **Agentic Coding Infrastructure:** Claude Code, Cursor, and GitHub Copilot integration for team development acceleration

LLM Security Engine (engsec_llm_engine):

- Self-correcting multi-agent system (Processor + Reviewer agents)
- RAG with OpenSearch vector store, Google Docs ingestion
- Jira prioritization automation (L1/L2 gates, risk analysis, P0/P1/P2 assignment)
- FastAPI service with M3 metrics, Flipr feature flags

Data Security & Privacy:

- **DataK9:** ML-based data classification and protection
- **Privacy Export:** DSAR automation with LangFx service
- **Privacy Guard:** Real-time privacy monitoring
- **GDPR Compliance Automation:** Regulatory compliance tools

Data Engineering Pipelines:

- Apache Spark (PySpark), Kafka streaming, Presto SQL
- ETL/ELT pipelines, incremental processing, backfill automation
- Piper workflow orchestration, YARN resource management

Technologies: Python 3.9/3.11, LangChain, LangGraph, FastMCP, PyTorch, TensorFlow, Apache Spark, Kafka, Presto, Pandas, OpenSearch, Pinecone, Qdrant, FastAPI, Protocol Buffers, Bazel, pytest

FasterOutcomes | Chief Technology Officer

June 2025 - Present | Remote

AI-powered legal technology startup automating legal workflows with production LLM systems. Started part-time, built MVP in 12 weeks, went full-time June 2025 to scale to Series A.

Leading technical strategy, product architecture, and engineering execution for AI legal tech platform serving law firms and corporate legal departments. Built production LLM applications from 0→1 with focus on RAG systems, AI agents, and agentic workflows.

Technical Leadership & Product Strategy:

- **Architected complete AI platform** from ground up: Production RAG system for legal document analysis, AI agents for contract review, agentic workflows for multi-step legal processes, vector database integration (Pinecone)
- **Led product roadmap** and go-to-market strategy: Identified product-market fit, designed user workflows, prioritized features

- **Built MVP in 12 weeks** with small team: 0 to production-ready platform with real customers and revenue
- **Technology stack decisions:** Backend (Python, FastAPI, LangChain, LangGraph, Firebase Functions), Frontend (React, NextJS, TypeScript, TailwindCSS), Cloud (GCP - Cloud Run, Firestore, GCP Storage), AI (OpenAI GPT-4, Anthropic Claude, Cohere)

Business & Team:

- Hands-on technical work (70%) + strategic leadership (30%)
- Hiring and mentoring engineering team
- Stakeholder management (investors, customers, legal advisors)

Key Achievements:

- Shipped production LLM features serving real law firms
- Achieved product-market fit in legal AI automation space
- Built scalable AI platform with modern cloud-native architecture
- Maintained security and privacy compliance for sensitive legal data

Technologies: Python (FastAPI, LangChain, LangGraph), React, NextJS, TypeScript, Firebase/Firestore, GCP (Cloud Run, Cloud Functions), OpenAI, Anthropic Claude, Pinecone, TailwindCSS, shadcn/ui

BayRockLabs (MatchPoint Solutions) | AI Solutions Lead & Technical Interviewer

January 2021 - December 2025 | Remote

GenAI consulting and technical recruiting for top tech companies

GenAI Lead for Client Projects:

- Led AI/ML implementations across industries: ecommerce, manufacturing, accounting, legal, hospitality
- Architected production LLM solutions (RAG systems, AI agents, chatbots)
- Delivered custom AI automation for client workflows
- Full-stack implementations (Python backend, React frontend, cloud deployment)

Technical Interviewing (100+ Interviews Conducted):

- Conducted technical interviews for **top tech companies:** Uber, Robinhood, Netflix, Airbnb, Peloton, Gap, and others
- Evaluated candidates for senior engineering roles (Staff, Principal, Senior)
- Interview domains: Backend, Frontend, System Design, AI/ML, Security
- Built deep network across tech industry through interview work

Notable Projects: Counta AI (agentic accounting automation), Lucid AI Solutions (ecommerce AI), Manufacturing AI (predictive maintenance), Legal Tech AI (document analysis)

Technologies: LangChain, LangGraph, Python, React, NextJS, Firebase, GCP, OpenAI, Anthropic, various client-specific stacks

Home Depot Canada | Lead Solution Architect

2017 - 2020 (3 years) | Toronto, ON

Led enterprise architecture and development for Home Depot Canada's digital transformation initiatives, focusing on ecommerce platform, mobile applications, and omnichannel customer experience.

Key Responsibilities:

- **Led mobile app development** (iOS/Android) for Home Depot Canada (Native iOS - Swift, SwiftUI; Android - Kotlin, Java; 1M+ downloads, 4.5+ star rating)
- **Architected ecommerce platform** on SAP Hybris (Microservices architecture with Spring Boot, integration with legacy systems)
- **Led team of 10+ engineers** across mobile, web, and backend
- **Technology stack modernization:** Java/Spring Boot, Angular, React, cloud migration

Key Achievements:

- Launched mobile app serving 1M+ Canadian customers
- Architected omnichannel platform (web, mobile, in-store integration)
- Led cloud migration from on-premise to hybrid cloud
- Reduced page load times by 40% through architecture optimization

Technologies: Java (Spring Boot), SAP Hybris, iOS (Swift), Android (Kotlin), Angular, React, Microservices, REST APIs, Azure, CI/CD

TD Bank I Senior Manager, Enterprise Applications

2014 - 2017 (3 years) | Toronto, ON

Led enterprise application development and architecture for TD Bank's internal systems, managing large-scale .NET applications and data platforms.

Key Achievements:

- Delivered enterprise-scale banking applications serving thousands of employees
- Led successful migration from legacy systems to modern .NET stack
- Implemented data warehouse for business intelligence and reporting
- Achieved 99.9% uptime for critical banking applications

Technologies: .NET (C#, ASP.NET), SQL Server, SSIS/SSRS/SSAS, Data Warehousing, Enterprise Architecture

Canada Life I Enterprise Architect

2012 - 2014 (2 years) | Toronto, ON

Key Achievements:

- Designed enterprise architecture framework using TOGAF 9.1
- Led development of reusable component libraries
- Architected multi-tenant insurance platform

Technologies: .NET, Angular, TOGAF, Enterprise Architecture, Microservices

CAA South Central Ontario I Senior .NET Developer

2010 - 2012 (2 years) | Toronto, ON

Key Responsibilities:

- Full-stack .NET development for CAA membership systems
- SAP Hybris ecommerce implementation
- Integration with legacy systems

Technologies: .NET (C#, ASP.NET), SAP Hybris, SQL Server, Web Services

Key Achievements & Awards

Uber:

- Multiple Internal Innovation Awards - ELLE AI Decision Engine (EngsecLLM Engine)
- Engineering Excellence Award - 96% test coverage on authorization middleware
- Automated 70% of privacy reviews - Reduced review time from days to minutes
- Built 9 enterprise web applications - Serving thousands of internal users
- 50+ security infrastructure projects - Across Python monorepo
- IAB TCF v2 Certification - Web CMP consent management platform

Technical Achievements:

- 96% code coverage on production authorization middleware (exceeding 85% standard)
- 13,000+ LinkedIn followers - Active thought leader in AI/ML and agentic coding
- Published technical articles on Medium (@FrankGoortani)

- 100+ technical interviews conducted for top tech companies (Uber, Netflix, Airbnb, Robinhood)
- Open source contributions - GitHub repositories with AI/ML tools and examples

Leadership & Impact:

- Led teams of 10+ engineers at Home Depot and TD Bank
- CTO - FasterOutcomes (AI legal tech startup, 0→1→Series A)
- Built products from 0→1 - MVP to production in 12 weeks
- Shipped production AI systems - LangChain, RAG, AI Agents at scale

Education

Master of Science (M.Sc.) - Management

AmirKabir University of Technology (Tehran Polytechnic), Tehran, Iran
Graduated: 2003

Bachelor of Science (B.Sc.) - Computer Software Engineering

AmirKabir University of Technology (Tehran Polytechnic), Tehran, Iran
Graduated: 2001

Professional Certifications

- TOGAF 9.1 Certified - The Open Group Architecture Framework (2013)
- PMP (Project Management Professional) - Project Management Institute (2013)
- Certified Data Scientist - Johns Hopkins University (2015)
- ITIL Foundation - IT Infrastructure Library (2013)
- Microsoft MCPD - Web Application Development (2012)
- Microsoft MCITP - Database Developer (2012)

Teaching & Speaking

College Instructor:

- George Brown College - Web Development and Programming
- Cestar College - Software Engineering
- Lampton College - Computer Science

Corporate Training:

- Learning Tree International - Technical instructor
- The Judge Group - Enterprise software training

Publications & Thought Leadership

Medium Blog: medium.com/@FrankGoortani

Technical articles on AI/ML, production LLM systems, and software architecture. Topics: RAG systems, LangChain/LangGraph, AI agents, Claude Code, agentic coding workflows, MCP development, distributed systems

LinkedIn Thought Leadership:

13,000+ followers with regular posts on agentic coding, Claude Code, MCP development, AI/ML engineering, production LLM systems. Active engagement with AI engineering community.

Languages

- **English:** Fluent (Professional Working Proficiency)
- **Farsi (Persian):** Native

Additional Information

Current Status:

- Full-Time CTO at FasterOutcomes (AI legal tech startup, June 2025-Present)
- Scaling company from early revenue toward Series A
- Leading engineering team growth and technical strategy
- Not currently seeking new opportunities

Work Preferences:

- 100% remote (non-negotiable for health reasons: back/shoulder issues)
- Async-first culture with minimal meetings
- Results-oriented, high-trust environment
- Modern tech stack (AI/ML, cloud-native, microservices)

Business Entities:

- VisionzLab - AI-driven architecture and development studio
- Architect.solutions - AI-enabled development consulting
- Focus: Agentic coding, Claude Code expertise, MCP development, production LLM systems

Aurora, Ontario L4G 6P7 (647) 983-5277 frank.goortani@gmail.com linkedin.com/in/frankgoortani github.com/FrankGoortani