Dmytro Komarchuk, PhD | Lead Software Engineer

San Francisco, CA • dmitruyk@gmail.com • +16287248570
LinkedIn https://www.linkedin.com/in/ikodi/

Technical Architect and Engineering Leader with extensive experience designing and scaling mission-critical distributed systems across AI, cloud infrastructure, fintech, and IoT domains. I lead architecture for complex, high-impact platforms that integrate advanced AI capabilities with production-grade reliability, security, and performance at scale.

Recognized for driving end-to-end technical strategy — from long-term architectural vision and platform standardization to execution across multi-team environments. I operate at the intersection of AI research, cloud engineering, and product delivery, transforming experimental innovation into scalable, enterprise-ready systems serving real-world operational needs.

Deep expertise in distributed systems design, multi-agent AI orchestration, high-throughput backend services, event-driven architectures, and cloud-native infrastructure. I emphasize strong engineering fundamentals: fault tolerance, observability, performance optimization, secure-by-design principles, and clean, evolvable system boundaries.

Experienced in influencing cross-functional stakeholders, mentoring senior engineers, defining technical roadmaps, and raising the engineering bar through architecture reviews, design standards, and scalable platform patterns.

Technical Depth:
Python · Go · C/C++ · Distributed Systems · Cloud-Native Architecture · Kubernetes · AWS · High-Scale APIs · PostgreSQL · Event-Driven Systems · OAuth/OIDC · Security Architecture · LLM & RAG Systems · Scalable AI Infrastructure

Key Skills:
Languages: Python, C/C++, Solidity, GoLang, SQL, bash/shell
Technologies: GQL, REST, Web Services, WebSockets, Docker, RabbitMQ, Kafka, Neo4j, Protobuf, Django, TCP/IP, HTTP, DNS, MQTT, FastAPI, Flask, TimescaleDB, ArgoCD, Prometheus, Linux, CI/CD, Shell, Retrieval-Augmented Generation (RAG), Finetuning, Amazon SageMaker, Strandsagents
Methodology: Scrum, Kanban, Waterfall, Teal company
Databases: PostgreSQL, Redis, TimescaleDB, MySQL, MongoDB, Vector DBs (e.g., FAISS, Weaviate, Qdrant)
Application Servers: Nginx, Apache *, unicorn, daphne, supervisord, GCloud, Spark
Cloud: AWS, GCP, Azure

Work Experience:

Senior Engineer, Cloud | BrightAI, *Palo Alto, CA — Sept 2025 – Present*

- **Summary***:* As a Cloud Engineer at BrightAI, I design and scale cloud-native backend systems powering AI-driven enterprise applications in industrial and physical AI domains. I lead architecture for multi-agent LLM systems, distributed microservices, and high-throughput APIs, ensuring production-grade reliability, security, and scalability. I collaborate closely with AI research, product, and platform teams to translate complex operational requirements into resilient, intelligent cloud solutions.
- **Main Duties & Responsibilities:**
  - ***AI Agent Architecture & Orchestration:***
    - Designed and implemented a production-ready multi-agent LLM pipeline with abstract BaseAgent, protocol-based envelopes/results, and strongly typed Pydantic DTO contracts.
    - Built shared AgentContext (LLM registry, RAG services, provider abstraction) enabling per-agent model selection (OpenAI / AWS Bedrock).
    - Architected full orchestration flow: Orchestrator → Routing Classifier → Data Extractor → Knowledge

Planner → Strategic Planner → Execution Coordinator → Final Response.
    - Implemented structured strategic planning with step dependencies, parallel execution, fallbacks, and traceable execution graphs.
    - Developed intelligent routing to dynamically adapt reasoning flows to industrial domain contexts.
  - ***Cloud Architecture & Distributed Systems***
    - Architected FastAPI-based microservices using async/await and Python concurrency best practices.
    - Designed scalable backend systems with PostgreSQL (SQLAlchemy 2.0), repository patterns, and clean architecture principles.
    - Implemented middleware for request tracing, per-user rate limiting, high-throughput concurrency control, and structured logging.
    - Optimized system latency and throughput for real-time AI inference workloads**.**
  - ***Retrieval-Augmented Generation (RAG) & Data Systems***
    - Built RAG pipelines with S3-backed document storage and embedding-based retrieval.
    - Designed knowledge tree construction and structured reasoning layers for controlled LLM outputs.
    - Managed schema migrations with Alembic and maintained versioned data contracts via Pydantic v2.
  - ***Enterprise Integrations & Security***
    - Integrated external AI services (Caterpillar AI Assistant, OpenAI) using OAuth 2.0 (Microsoft Entra ID) with secure token auto-refresh.
    - Implemented unified authentication supporting JWT (RS256/HS256), OIDC (Auth0), and Basic fallback for administrative access.
    - Managed typed configuration with environment precedence and SOPS-encrypted secrets mounted via Kubernetes.
    - Ensured compliance-ready security posture across services**.**
  - **DevOps, CI/CD & Platform Reliability**
    - Deployed containerized services using Docker and Kubernetes.
    - Contributed to infrastructure-as-code workflows and cloud-native operational practices.
    - Implemented observability stack: Sentry tracing, health checks, async execution tracing, and structured monitoring.
    - Improved deployment safety, rollback strategies, and runtime visibility.

- **Success Characteristics:**
  - Demonstrated technical ownership of core AI orchestration architecture.
  - Improved system scalability and fault tolerance through asynchronous design and distributed service boundaries.
  - Reduced operational risk via strong typing, schema validation, and trace-based debugging.
  - Mentored engineers on async programming, cloud-native architecture, observability, and secure API design.
  - Bridged AI research and production engineering, translating experimental models into reliable enterprise systems.

- **Working Conditions & Requirements:**
  - Operate in a fast-paced, AI-first product environment with iterative releases.
  - Participate in cross-functional design reviews, roadmap planning, and architecture discussions.
  - Manage concurrent initiatives across AI platform, cloud infrastructure, and enterprise integrations.
  - Collaborate with distributed teams across engineering, AI research, and product.

Senior Staff Software Engineer | Level Home Inc, *Aug 2022 – Present*

- **Summary:** As an Applications Engineer, focused on developing core subsystems, choosing the right tools and technologies, and ensuring the scalability, reliability, and simplicity of the platform. Worked on an ambitious roadmap, handling tasks ranging from cryptographic operations to various IoT protocols like WiFi, Ethernet z-wave, bluetooth.
- **Main Duties & Responsibilities:**
  - Architected, designed, and implemented backend services including the core API and backend data processing.

- Owned the full lifecycle of development including design, prototyping, testing, release, and deployment.
- Programmed primarily in Python, Go and C, and interacted directly with other engineering and product teams.
- Automated testing of contributed code and conducted daily code reviews.

- **Success Characteristics:**
  - Demonstrated ownership by taking features from requirements through to bug fixes.
  - Provided constructive feedback during the code review process as a mentor and adapted to changing requirements and priorities in an agile environment.
  - Kept up with technology trends to improve code quality, patterns, and processes.
  - Advocated for process improvements with suggestions.
- **Working Conditions & Requirements:**
  - Regular attendance and active participation in daily stand-up meetings.
  - Handled job-related stressors effectively and maintained unimpaired performance.
  - Occasionally traveled for business purposes.

Senior Software Engineer DataArt, June 2021 – Aug 2022

- **Summary:** As a Senior Software Engineer, focused on developing a finance risk prediction system based on statistical data, using Linux, Python, AI/ML technologies. Worked on creating a web application for estimating credit and liability parameters of funds over different investment periods. Ensured high performance, reliability, and scalability of the system.
- **Main Duties & Responsibilities:**
  - System Design & Implementation: Led core component design and development, ensuring robust architecture and scalability.
  - Data Preparation: Managed data preprocessing, including cleaning, normalization, and feature engineering for ML analysis.
  - AI/ML Integration: Integrated machine learning models, selected algorithms, tuned hyperparameters, and validated models to enhance accuracy.
  - Parallel Processing: Implemented parallel computation to boost data processing and model training efficiency, reducing runtime.
  - Cloud Functions: Deployed AWS Lambda and Google Cloud Functions for scalable data processing and model inference.
  - Performance Optimization: Improved algorithms and data pipelines, profiling and refactoring code to enhance performance.
  - Collaboration: Coordinated with data scientists, analysts, and engineers to align technical solutions with business needs.
  - Automated Testing & QA: Automated testing processes, including unit, integration, and performance tests, to ensure application reliability.
  - Documentation: Developed documentation for systems, algorithms, and processes to aid maintainability and knowledge sharing.
  - Mentorship: Guided junior engineers, promoting a collaborative and learning-focused environment.

Principal Software Engineer Autility, *Nov 2019 – June 2021*

- **Summary:** Led the design and implementation of a digital twin system for building operations, focusing on remote control and efficiency improvements. Drove architectural decisions, managed a team of 25, and collaborated with stakeholders to deliver innovative solutions.
- **Main Duties & Responsibilities:**
  - Architected and implemented a digital twin system to control and optimize building operations, significantly enhancing operational efficiency.
  - Owned the full lifecycle of development, from initial design and prototyping to deployment and maintenance.
  - Programmed mainly in Linux, bash/shell, Python and C, with a focus on scalable and reliable solutions.
  - Worked closely with product owners, UX/UI designers, and other engineering teams to align technical solutions with business goals.
  - Automated testing processes and conducted regular code reviews to ensure high-quality code.

- **Success Characteristics:**
  - Demonstrated strong ownership by leading projects from requirements gathering to deployment and maintenance.
  - Acted as a mentor by providing constructive feedback during code reviews and fostering a collaborative team environment.
  - Adapted to changing requirements and priorities, maintaining agility in the development process.
  - Proactively brought in new technologies and best practices to improve code quality and development efficiency.
  - Advocated for continuous improvement in processes and workflows.

Senior Software Engineer VLC Energy GmbH, *Sep 2018 – Nov 2019*

- Developed a full-stack billing system for electric vehicle charging stations using Linux, Docker, Python, Django, C, and Bash, supporting over 10,000 daily sessions and integrating a network of chargers via Ethernet and Wi-Fi protocols. Led end-to-end development of mobile, web, and backend components, completing the project ahead of schedule and within budget.

Scientist, Researcher, Lecturer NULES UNIVERSITY, *Sep 2014 – Aug 2018*

- Developed software for crop status estimation using UAV imagery using Linux, bash/shell, C++ and Python, optimizing fertilizer use and increasing crop yields by up to 50%.

Education

Ph.D. in Electrical Systems (Hardware and Software)
National University of Life and Environmental Sciences of Ukraine
2010 – 2013

- Focus: Electrical Complexes and Robotic Systems
- Dissertation: *Mathematical Modeling of Induction Heaters and Development of Automation Systems*

Bachelor's Degree in Economics and Entrepreneurship
National University of Life and Environmental Sciences of Ukraine
2009 - 2010

Awards: Laureate of the President Prize for Young Scientists 2018