

Raman Marozau



Warsaw, Poland | +48 881 592 968 | raman.research@worktif.com

[LinkedIn](#) | [Website](#) | [Runtime Web Framework](#) | [studio Bus Kernel](#)

13+ years building and **scaling backend** and **cloud systems**, **leading teams**, and running delivery in **fintech**, **banking** and **high-load environments**. I design serverless and event-driven architectures on AWS, automate infrastructure with CDK/SDK and Unix/C, and keep systems observable and safe to evolve.

I focus on: turning vague business requirements into concrete, running systems; removing manual work from infrastructure; and pushing performance (from Lambda cold starts to low-level tensor kernels).

Selected Results

- **AI performance:** Implemented custom low-level tensor kernels (SIMD, FP16/AVX512) and achieved 20-30x speed-up for key deep-learning operations in production-oriented benchmarks.
Serverless banking systems: Designed and rolled out AWS Lambda-based architectures for fintech/banking workloads, so new features could be shipped without touching “core” monoliths and with predictable cloud costs.
- **Infra automation:** Replaced manual console work with CDK-driven IaC and Unix/C automation pipelines, so new environments and CI/CD flows are brought up by code, not human procedures.
- **Team impact:** Led cross-functional teams up to ~14 people and owned delivery end-to-end (discovery → scope → implementation → rollout → support) on large financial-services projects.
- **Architecture risk insight:** Built an engine that continuously analyses microservice dependencies/configs and surfaces architectural risks without manual system mapping.

Skills (Condensed)

- **Backend & Serverless:** TypeScript/Node.js, async I/O, functional/OOP design, event-driven microservices, REST/GraphQL; AWS Lambda, Layers, API Gateway, EventBridge, S3, CloudFront, IAM; PostgreSQL, RDS, NoSQL, caching, robust error-handling and retries.
 - **Infrastructure, IaC & DevOps:** Unix/Linux, networking, DNS, HTTP, reverse proxies/load-balancing; AWS CDK, CloudFormation; CI/CD, Git workflows, blue-green/canary, Docker; observability (logs, metrics, tracing); security (least-privilege IAM, secrets, encryption, scanning, policy-as-code).
 - **Frontend & Product:** React (primary), Angular, Material UI; Redux + RxJS, React Context, complex forms; SPA architecture, performance tuning (bundling, lazy loading, SSR/CSR decisions), UX-oriented flows and error-handling.
 - **AI & Performance:** Deep-learning fundamentals, LLM adaptation and evaluation; tensor-operation optimisation (SIMD/AVX2/AVX512/NEON; FP32/FP16/BF16); serverless AI deployment on AWS; cold-start mitigation, caching; governance, data ownership, safety-by-design, auditability.
 - **Leadership:** End-to-end delivery ownership, cross-functional coordination, technical decision-making, mentoring, expectation management, and keeping work tied to business outcomes instead of just ticket closure.
-

Professional Experience

Cloud Architecture & Automation Consultant

Warsaw, Poland – Apr 2025 – Sep 2025

- Architected and deployed AWS infrastructures with CDK and SDK so that all environments are created and updated via code instead of manual console changes.
- Built Unix-based automation (shell + compiled C modules) for infra bootstrapping, monitoring and CI/CD, removing ad-hoc scripts and snowflake setups.
- Took ownership of complex cloud programs from “idea” to production, coordinating backend, cloud and DevOps work around a single delivery plan.

- Audited existing cloud usage and restructured architectures to cut infrastructure spend while keeping elasticity and reliability (function-level controls, usage-based patterns).
-

Team Leader – Armada Labs

Warsaw, Poland – Jun 2022 – Apr 2025

- Built and led a **cross-functional team up to 14 people** (backend, frontend, QA, analysts, managers) on fintech and banking-infrastructure projects; set cadence, priorities and removed blockers before they hit delivery.
- Worked directly with banking stakeholders: unpacked requirements, mapped them to concrete features and internal process changes, and tied all of it to regulatory and risk constraints.
- Owned large-scale initiatives end-to-end: led discovery workshops, defined MVP vs. later phases, and aligned timelines/budget with what engineering could actually deliver.

Designed serverless and event-driven architectures on AWS (Lambda microservices, event orchestration, IaC) so that new financial products and internal tools could be delivered without growing ops overhead.

- Explained architectural decisions in business language (risk reduction, compliance fit, TCO), which kept stakeholders aligned and reduced “architecture vs. business” conflicts.
-

Management Staff & Team Leader – Vention

Remote – Sep 2020 – Jun 2022

- Led and coordinated an engineering team: set clear delivery goals, prioritised work, and kept focus on outcomes (what has to run in production by when), not on ticket counts.

- Managed hiring and onboarding: selected engineers who fit both the stack and the way the team worked; helped them grow into senior roles through direct mentorship and feedback.
 - Turned cross-department and client requests into executable delivery plans: clarified scope, negotiated priorities and kept projects on track against time and budget.
 - Played hands-on lead full-stack role: designed and shipped features across the stack, reviewed critical code paths and steered architecture towards maintainable, observable systems.
-

Lead Software Engineer & Team Leader – Softeq Development Corp.

Remote – Jan 2017 – Jun 2020

- Delivered features end-to-end across frontend and backend: clarified requirements, designed solutions, implemented, reviewed, deployed and stabilised them in production.
- Led a team of engineers in a dev+team-lead model: split work, aligned on approach, and created a space where engineers could challenge design decisions instead of just “taking tasks”.
- Kept project scope, milestones and deliverables aligned with product goals by working through project/account managers and translating customer feedback into technical backlogs.
- Found and fixed performance bottlenecks across DB, backend services and frontend: optimised queries, refactored hot paths, improved rendering/interaction, and measured impact via latency and UX.

Junior → Middle → Senior Software Engineer (Earlier Roles)

ASW Systems a.s., Createch Group, Piplos Media, New Cite – 2012 – 2017

- Implemented custom React + RxJS state-management flows for complex UIs and multi-step forms with strict validation in a HORECA (restaurant/POS) information system.

- Built business charting and analytics UIs in React, tightly integrated with server-side data, and improved client–server exchange to keep dashboards responsive.
 - Delivered social-media and digital-selling platforms (Symfony2 + modern CMS), gaining experience in user-behaviour analysis and CRM-style systems.
 - Worked across backend, DB and Symfony2 in a large digital community; delivered an external management portal for marketing data sampling in e-commerce.
-

Entrepreneurial Experience

Founder & Principal Architect (CEO) – Work Target Insight Function

Warsaw, Poland – Apr 2024 – Present

- Built Lambda Kata – deterministic Node.js runtime for AWS Lambda that cuts cold-start jitter and tail latency for serverless workloads and makes behaviour reproducible across environments.
- Designed “stdio Bus kernel” – low-level C protocol and runtime for agent systems that routes messages over stdin/stdout and Unix sockets, with explicit backpressure and queue control for AI agents and tools.
- Built hardware-accelerated ML components – custom tensor kernels (SIMD, FP16/AVX512) that give 20-30x acceleration for selected deep-learning operations.
- Created Runtime Web – full-stack TypeScript/Node.js framework on AWS (CDK/SST) that can spin up production-ready microservice backends and AI chat endpoints from a minimal config instead of manual wiring.
- Integrated infra, backend and frontend workflows into a single pipeline, so performance optimisations (from C/kernel level up to React) are applied consistently, not as isolated hacks.
- Applied ideas from quantum-physics and numerical research to concrete optimisation patterns in code, focusing on predictable latency and efficient resource usage for AI-driven systems.

Founder & Principal Architect (CEO) – Machine27

Warsaw, Poland – Feb 2025 – Present

- Built an architecture and risk engine that continuously scans service dependencies, data flows and configs to surface structural weaknesses in complex systems.
 - Automated system-architecture analysis: the engine generates recommendations without manual diagramming and keeps a live view of risk as systems evolve.
 - Wired risk detection and change-planning into CI/CD, so architectural drift is caught alongside tests and deployments, not months later.
-

Education

MSc, Technical Sciences – Engineering Psychology & Ergonomics

BSUIR, 2013 – 2015 – Thesis: smart-home system integrating multiple electronic subsystems into a unified, software-controlled environment with safety and performance guarantees.

BSc, Radioengineering Systems

BSUIR, 2008 – 2013 – Classical radioelectronics and radioengineering: electromagnetic theory, antennas, broadcasting/navigation systems, circuit design and embedded/software control of hardware.

Research

- **AGI & AI Optimisation** – Designed and analysed distributed AI/multi-agent architectures that decompose complex tasks into cooperating agents, with explicit constraints on safety, interpretability and failure modes. Built prototype systems that route work between agents and measure how coordination patterns affect solution quality and latency.

- **Quantum Decoherence & Entanglement Dynamics** – Developed conceptual and mathematical models of information propagation and stabilisation in closed quantum-like systems, focusing on internal decoherence, effective stress–energy components and their impact on early-universe–style dynamics. Authored manuscripts exploring how these mechanisms can inspire more robust AI architectures.
- **Mathematical Adaptation of LLMs** – Ran experimental benchmarks of tensor layouts, numeric precision (FP32/FP16/BF16) and custom kernel implementations on latency, throughput and numerical stability for production-grade LLM usage. Implemented and profiled low-level kernels (SIMD/AVX/NEON) to quantify speed–accuracy trade-offs in real workloads.
- **Security, Data Ownership & AI Governance** – Formulated the Open Data Ownership License concept and designed governance patterns that tie AI model usage to explicit data-ownership and accountability rules. Worked out enforcement flows for access control, auditability and revocation in AI-driven systems.

Languages

- English: Advanced
- Belarussian: Fluent (native)