ALEX KOSACHOV

Principal Data Architect · Strategic, Cross-Cutting Data Architecture for Regulated Multi-Domain Platforms

Canadian PR · Open to Ontario relocation · Fredericton, NB
+1 506 282 3408 | alex@kosachov.ca | linkedin.com/in/kosachov | github.com/wastrox

Summary

- 17 years of experience in engineering leadership at SCG Global — Senior Developer (2009), Director of Engineering (2014), Chief Technology Officer (2019 to present); 12 years at director level and above, 7 years as CTO.
- 10+ years of experience defining and operating strategic, cross-cutting data architecture across a regulated, multi-domain enterprise platform (8 integrated products, 70+ repositories, 27 vessels, 82,666 active records, $226.7M reconciled).
- Proficient in canonical data modeling, schema standardization, data contracts, multi-layer persistence strategy, event-driven architecture (Kafka Exchange), and batch / micro-batch / streaming pattern selection.
- Excellent skills in cost optimization, workload rationalization, storage tiering, partitioning, indexing, deduplication, and technology evaluation for modern data stacks (Databricks vs Snowflake, Delta vs Iceberg vs Hudi, lakehouse vs warehouse).
- Production experience with AI and GenAI enablement: OpenAI API, prompt engineering, human-in-the-loop (HITL) validation, agentic automation pipelines, and team adoption of GitHub Copilot and Claude Code; active ramp on RAG pipelines, vector stores, and embeddings (April 2026).
- Strong regulated-industry background: Bureau Veritas classification certification (personally led the audit), legal-grade audit trail design, 92 production releases with zero data loss and zero Port State Control detentions; architectural familiarity with HIPAA, FHIR, USCDI, and PHI / PII governance.
- Worked with product leadership, executive stakeholders, and external auditors across all phases of the SDLC; built and scaled an engineering organization of up to 12; completed a three-year succession arc to resolve the single-threaded-custodian pattern.

Tools

- **Languages & Frameworks:** Ruby on Rails, JavaScript, React, Node.js
- **Databases:** MySQL, MS SQL
- **Cloud & Infrastructure:** AWS (EC2, wartime on-prem to cloud migration), Docker, Nginx, Hetzner
- **Event-Driven & Integration:** Kafka, RabbitMQ, SymmetricDS, REST / GraphQL APIs
- **Observability & Monitoring:** Prometheus, Elasticsearch, Logstash, Grafana, Slack / Google Chat alerting
- **AI / GenAI (production):** OpenAI API, production prompt engineering, Human-in-the-Loop (HITL) validation, agentic automation pipelines, GitHub Copilot, Claude Code
- **Modern Data Stack (architectural familiarity):** Databricks, Snowflake, Delta Lake, Apache Iceberg, Hudi, Lakehouse vs Warehouse, OLAP, Vector Stores, Feature Stores, Semantic Layers
- **AI-Data Patterns (architectural familiarity):** RAG Pipeline Architecture, GenAI Foundations, LLM Data Preparation, Feature Engineering
- **Healthcare Standards (architectural familiarity):** FHIR, USCDI, HIPAA, PHI / PII Governance, Interoperability Standards
- **Active Ramp (April 2026):** Hands-on RAG pipeline proof (in progress), vector store + embeddings production patterns, healthcare-specific vocabulary (30-day ramp)

Experience

Chief Technology Officer **05/2019 to present**
SCG Global https://www.scglobal.group/ **Cyprus / Bulgaria / Ukraine (Full-time Remote)**

Strategic, cross-cutting engineering and data architecture ownership for an 8-product integrated maritime enterprise platform: 70+ repositories, 27 vessels, 82,666 seafarers tracked, $226.7M cumulative fleet operations, 1,000+ users across vessels and shore offices. A 12-engineer team delivered scope comparable to competitor platforms running 40+.

Cross-Product Data Coherence and Schema Standardization

- Defined and drove architecture standards across 8 products to prevent architectural drift: schema naming conventions, event patterns, canonical data models, data representation standards — replaced fragmented local conventions with platform-wide discipline that held through three team rotations and 92 releases.
- Architected event-driven architecture (EDA) via a Kafka-based Exchange decoupling 8 products with no direct database coupling; applied batch, micro-batch, and streaming pattern selection per domain based on access patterns and consistency requirements; zero cross-product cascade failures across 92 production releases.
- Designed multi-layer persistence strategy on a production OLTP base (MySQL, MS SQL) with multi-database per-tenant sovereignty (27 vessel instances + office replica) — clean data jurisdiction per tenant, zero cross-contamination across ownership transitions and decommissioning events.
- Established a data contract codified as a 15-year backward-compatible API across fleet running heterogeneous software versions simultaneously — 316 database migrations, zero failures.

- Owned the canonical data model across the 8-product portfolio: Shipvisor (PMS), Finvisor (finance, $226.7M reconciled), Crewisor (crew, 82,666 records), Invisor (procurement), E-SCMS (compliance), Heatech (analytics), DaX (emissions), AI-powered recruitment.

Cost Optimization and Operational Efficiency

- Delivered 17-18x fleet-wide query performance improvement through targeted partitioning, indexing strategy, selective denormalization, and caching discipline — no rewrite, no compute spend increase, no new licensing.
- Deduplication at production scale: Jaro-Winkler fuzzy-match pipeline for seafarer duplicate detection across the 75K+ candidate database; article-identity deduplication at 67.8x compression factor across 1.1M placements in the parts catalog.
- Workload rationalization through architectural leverage: absorbed 60%+ fleet growth without shore office headcount increase — platform discipline, not staff addition.
- Designed and operated an observability stack from first principles (Prometheus, Elasticsearch, Logstash, Grafana, Slack / Google Chat alerting) — open-source alternative to commercial APM, eliminated six-figure annual licensing; informed by OpenTelemetry's three-pillar model.
- Led technology evaluations across the stack based on cost, performance, governance, and operational complexity — selected SymmetricDS over commercial alternatives for real-time replication; chose Kafka over proprietary messaging; selected MySQL over Postgres for specific workload shapes. The evaluation muscle transfers directly to Databricks vs Snowflake, Delta vs Iceberg vs Hudi, and lakehouse vs warehouse decisions the PDA role drives.

AI and GenAI Enablement

- Architected data foundations for AI production workloads: OpenAI API crew recruitment automation — LLM-powered CV extraction, candidate matching against 75K+ seafarer database, task routing with human-in-the-loop review before any database commit; 2,056 CVs processed, 20,087 sea-service entries extracted.
- Production prompt engineering: several complex prompts drive CV parsing and candidate-to-vacancy mapping under HITL validation.
- Directed team adoption of AI-assisted developer tooling (GitHub Copilot, Claude Code) across code review, documentation, and CI/CD — approach agreed through team consensus, implementation executed under architectural direction.
- Use AI as a catalyst for architectural productivity — accelerated pattern review, exploratory prototyping, documentation generation; mirrors PointClickCare's human-first, AI-accelerated operating philosophy.
- Active ramp: RAG pipeline architecture, vector stores, and embeddings — hands-on proof in progress (pet project, April 2026).

Regulated-Industry Architecture

- Personally led the Bureau Veritas classification certification audit — system passed, certification earned for commercial fleet operations.
- Designed legal-grade audit trail infrastructure across 8 products (1.4M+ financial audit entries, 2.9M PaperTrail crew-operation events) years before regulators asked — ISM Code compliance across 17 vessels for 8+ years, zero PSC detentions, 576 regulatory certificates tracked with zero gaps.
- Designed AES-256-CBC cross-runtime encryption (crypto-js to Ruby OpenSSL) before legal review required it — zero security incidents across 11.2 years of Heatech operation.
- Maritime regulatory discipline maps directly to HIPAA, FHIR, USCDI, and PHI / PII governance — pattern language already transfers; healthcare-specific vocabulary is a 30-day ramp.

Delivery, SRE, and Executive Partnership

- 92 production releases, 0 data loss, 0 PSC detentions across 15 years of continuous fleet operation.
- Led the 2022 wartime AWS infrastructure relocation from on-premise Ukraine — Starlink across the team, redundant connectivity, phased migration; 100% operational continuity, zero client loss, zero data loss.
- Bridged product and engineering toward a common data language — partnered directly with Head of Operations to translate fleet requirements into multi-year technical roadmap across 8 product domains.
- Balanced pragmatic urgency with long-term architectural integrity — shipped the backward-compat API contract in v1 when the fleet was 6 vessels, not 27; the discipline prevented architectural drift across 15 years, three team rotations, and 8 product additions.
- Built an engineering organization up to 12 across frontend, backend, QA, DevOps, DBA, sysadmin; ~15 hires over tenure; resolved the single-threaded-custodian pattern across a three-year succession arc — platform now operates without me as a bottleneck.

Director of Engineering **2014 to 2019**
SCG Global **Odesa, Ukraine**

- Led engineering organization through platform expansion from single-product PMS to multi-domain operational ecosystem — established foundational architecture standards, cross-module integration patterns, and data flow architecture that became the platform backbone.
- Built and scaled the engineering team; established code review practices, release governance, and delivery processes that still govern the platform today.

Senior Software Engineer / Lead Developer **2009 to 2014**
SCG Global **Odesa, Ukraine**

- Designed core backend architecture and data models for the Shipvisor maritime enterprise platform (Ruby on Rails + MySQL).
- Led a small development team through the initial platform modules: vessel maintenance, equipment lifecycle, document management, operational reporting.

Education

Master's Degree: Computer Engineering (Verified by WES Canada.)

Odesa I.I. Mechnikov National University, 1998 to 2003

Certifications & Professional Development

- Defending and Deploying AI — Pearson
- Certified Trainer, Personal Productivity & Leadership — FranklinCovey
- Advanced English, Conversation Skills — London School of English