

KARSTEN S. NIELSEN

781.810.4558 | Waltham, MA | karsten@skyt.com | www.linkedin.com/in/karsten-s-nielsen

SOFTWARE ARCHITECT

Software Architect | Strategic Technical Leader | Data Platform & Lakehouse Specialist

Pragmatic Software Architect who scales from guiding enterprise legacy modernizations to hands-on, solo execution of end-to-end data platforms. Expert in balancing technical rigor with business reality, delivering measurable impacts such as six-figure annual cost reductions and zero-downtime migrations. Bilingual (English/Danish) leader who excels at bridging the gap between executive stakeholders and distributed engineering teams, fostering collaboration to solve difficult architectural challenges.

CORE COMPETENCIES

Strategic Architecture: Legacy Modernization | Multi-Tenant Systems | Cloud Infrastructure | Cost Optimization | Regulatory Compliance (HIPAA, FERPA) | Lakehouse Architecture

Technical Leadership: Cross-Functional Team Leadership | Offshore Coordination | Mentorship | Agile Practices

Problem Solving: Root Cause Analysis | Performance Tuning | Zero-Downtime Migration | Requirements Analysis

TECHNICAL PROFICIENCIES

Architecture & Cloud: Microservices | SOA | Databricks Lakebase | Delta Lake (Medallion Architecture) | Containerization (Docker) | AWS (EC2, S3, VPC, IAM, Lambda) | Serverless Compute | C4 Model (Structurizr DSL)

Integration & Messaging: Event-Driven Architecture | RESTful APIs | SOAP | RabbitMQ | API Gateway | dbt-databricks | Healthcare Standards (EDI X12, HL7) | XML/XSLT

Data & Search: Elasticsearch (Expert: v5-8) | SQL Server | PostgreSQL 17 (Lakebase) | pgvector | Delta Lake | T-SQL | Data Modeling | NoSQL | Entity Framework

Languages & Frameworks: C# | JavaScript | .NET Core/6.0 | Python | Streamlit | C++ | Legacy (VB6, Classic ASP)

AI & Data Science: Retrieval-Augmented Generation (RAG) | LLM Agents (LangChain) | Vector Databases | XGBoost | Scikit-learn | Accelerated Data Science (RAPIDS/CuPy) | Deep Learning Fundamentals

DevOps & Quality: Terraform | GitHub Actions CI/CD | IAM OIDC | OAuth M2M | ruff | pyright | sqlfluff | pre-commit

PROFESSIONAL EXPERIENCE

SOFTWARE ARCHITECT

Ellevation Education | Boston, MA

2016 – 2025

Lead Architect for a prominent EdTech platform, responsible for modernizing critical infrastructure serving 1,500+ client databases in adherence with FERPA privacy standards.

- **Strategic Cost Optimization:** Identified inefficiencies in hosting architecture and designed a dynamic load-balancing algorithm based on real-time statistics. **Result:** Reduced hosting costs by **50% (\$100k+ annual savings)**.
- **Pragmatic Modernization:** Led the complex migration from .NET Framework to .NET 6.0. Avoided risky "big bang" releases by engineering a dual-framework operation path, allowing for gradual, safe modernization and containerization.
- **High-Volume Data Solutions:** Solved performance bottlenecks in a system processing 40+ tables across 1,500+ databases. Replaced underperforming ORM code with optimized stored procedures and intelligent hash-signature change tracking, reducing worst-case processing time from over a day to minutes.

- **Resilient System Design:** Architected an extensible indexing framework that allowed concurrent operation of multiple Elasticsearch cluster versions (v5 through v8). **Result:** Achieved a multi-year upgrade cycle with **zero downtime** for end-users.

SOFTWARE ARCHITECT

NaviNet | Boston, MA

2014 - 2016

Designed enterprise strategies for a healthcare transaction platform, establishing frameworks for geographically distributed teams.

- **Adaptable Design Implementation:** Translated complex healthcare compliance regulations (HIPAA) into configurable, scalable product solutions for Referral and Authorization workflows.
- **Global Collaboration:** Bridged the gap between local and offshore teams. Defined architectural guidelines that maintained consistency while allowing flexibility.
- **Hands-On Leadership:** Provided Proof of Concepts (POCs) for complex architectural patterns to guide implementation teams and conducted high-level code reviews to ensure alignment with design vision.

OPEN-SOURCE PROJECTS

SERVERLESS ANALYTICS PLATFORM

Open-Source (Apache-2.0) | AWS, Databricks

2026

End-to-end soccer analytics platform on Databricks Lakebase with medallion architecture, serving interactive dashboards from serverless PostgreSQL 17.

- **Infrastructure as Code:** Terraform (10 modules), IAM OIDC, OAuth M2M, KMS-encrypted state.
- **Data Engineering:** 30 dbt models with 256+ data quality tests across 5 providers, processing event streams and 25fps optical tracking data.
- **Applied Analytics:** Physics-based pitch control (Spearman 2017), expected threat (Markov chains), SPADL/VAEP action valuation (XGBoost).
- **Engineering Discipline:** Daily merge-to-main across 13 phases; 319 unit tests, 3 CI pipelines, automated pre-commit quality gates on every commit.

AUTOMATED ARCHITECTURAL GOVERNANCE — CLAUDE CODE PLUGINS

Open-Source (MIT)

2026

Reusable AI-assisted engineering plugins that enforce enterprise standards for architecture documentation, security auditing, and code quality. github.com/karsten-s-nielsen/mad-scientist-skills

- **Pre-Commit Quality Gate:** Automates code reviews, documentation checks, and C4 diagram regeneration (Structurizr DSL) on every commit — enforcing architectural governance.
- **Embedded Security Audit:** STRIDE threat modeling and OWASP Top 10 scanning with infrastructure hardening checklists, integrated directly into the development workflow.

PREVIOUS RELEVANT EXPERIENCE

NaviNet | Senior / Principal Software Engineer: Backend development for EDI transactions (270/271, 276/277)

Tieto | Principal Consultant: Clinical database systems for quality data tracking

Diamond Bullet Design | Consultant: C++ and early .NET consulting including an early E-commerce platform.

CERTIFICATIONS & TRAINING

Twelve Football: Soccermatics Pro (Enrolled – August 2026). Focus: Predictive analytics & tracking data (Python).

NVIDIA: Accelerated Data Science | Deep Learning | Building RAG Agents with LLMs | Adversarial ML (In Progress)

Leadership & Coaching: UEFA C License | Mental Performance Training

EDUCATION

Ph.D. Coursework (All Qualifying Exams Passed), Computer Science: University of Michigan | Ann Arbor, MI

Master of Science, Computer Science (Fulbright Scholar): University of Florida | Gainesville, FL

Bachelor of Science, Computer Science/Mathematics: Aarhus University | Aarhus, Denmark