

KARSTEN S. NIELSEN

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

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

Pragmatic Software Architect bridging formal academic rigor in cognitive AI with hands-on execution of modern, end-to-end data platforms. Trained in predictive human-computer interaction at the UMich AI Lab, I balance theoretical depth with business reality to deliver measurable impacts—from six-figure cost optimizations to AI-accelerated architectural governance. As a bilingual (English/Danish) leader, I excel at connecting executive stakeholders, distributed engineering teams, and data science collaborators.

CORE COMPETENCIES

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

Technical Leadership: Cross-Functional Team Leadership | Automated Architectural Governance | Offshore Coordination | Mentorship

Systems Engineering: Root Cause Analysis | Performance & Query Tuning | Zero-Downtime Migrations | High-Availability Design

TECHNICAL PROFICIENCIES

Architecture & Cloud: Databricks (Unity Catalog, Lakebase) | AWS (EC2, S3, VPC, IAM, Lambda) | Serverless Compute | Microservices | C4 Model (Structurizr DSL)

Data Engineering & Search: Delta Lake (Medallion Architecture) | dbt-core & dbt-databricks | Elasticsearch (Expert: v5-8) | PostgreSQL 17 | SQL Server | NoSQL

AI, ML & Analytics: Retrieval-Augmented Generation (RAG) | LLM Agents (LangChain) | Vector Databases (pgvector) | XGBoost | Scikit-learn | RAPIDS/CuPy

DevOps & Security: Terraform (IaC) | GitHub Actions (CI/CD) | IAM OIDC & OAuth M2M | ruff, pyright, pre-commit | Containerization (Docker)

Languages & Frameworks: C# (.NET Core/6.0) | Python | Streamlit (Databricks Apps) | JavaScript | C++ | T-SQL

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 migration from .NET Framework to .NET 6.0. Engineered a dual-framework path to avoid "big bang" releases, enabling 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 global 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

(RIGHT! LUXURY!) LAKEHOUSE — SERVERLESS SOCCER ANALYTICS PLATFORM

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

2026

- **Architecture & Scale:** Designed an end-to-end serverless soccer analytics platform on Databricks (Unity Catalog, Medallion architecture) processing 38M+ rows of high-frequency 25fps optical tracking data.
- **Infrastructure & Security:** Architected a zero-secrets CI/CD pipeline using GitHub Actions, AWS IAM, and Databricks OIDC federation, deploying 10 reusable Terraform modules with KMS-encrypted state.
- **Data Engineering:** Engineered robust transformation pipelines utilizing dbt-databricks (36 models, 381 data tests) to seamlessly normalize disparate coordinate systems from 5 professional data providers.
- **AI Models & Serving:** Built Streamlit dashboards natively on Databricks Apps, served by an autoscaling PostgreSQL 17 (Lakebase) backend leveraging pgvector for ML-driven player embeddings (Doc2Vec) and defensive valuations (XGBoost).

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.
- **Serverless Optimization Audit:** Automated profiling of data pipelines to identify algorithm bottlenecks, inefficient database queries, and cloud compute waste across distributed execution environments.
- **Cognitive Interface Audit:** Evaluates UI and developer workflows using predictive HCI frameworks (GOMS task models, NASA-TLX, Wood's error tolerance, Gergle visual grounding, accessibility).

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 (ABD), Computer Science: University of Michigan AI Lab | Ann Arbor, MI

Focus: Cognitive Architecture & Predictive HCI (GOMS, EPIC) under David Kieras & David E. Meyer

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

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