
Software engineer and architect with over 16 years of experience, specialising in data engineering, solution architecture, and technical leadership. Has scaled a college-founded platform to 100M+ users, cut pipeline failure rates by 90% within Google Cloud Storage, and turned 25 years of fragmented data into a single source of truth for a private equity client, designing and delivering data platforms, distributed systems, and AI-powered applications along the way. Brings both hands-on engineering depth and consulting fluency, translating ambiguous business problems into concrete technical proposals and driving them through to delivery. Equally comfortable in front of a client and in a codebase.

SKILL SET

LANGUAGES & FRAMEWORKS

Python, PySpark, Java, Spring Framework, TypeScript

AI & ML

RAG pipelines, LlamaIndex, LangChain, ChromaDB, Gemini API, vector databases

CLOUD & INFRASTRUCTURE

AWS, GCP, Docker, Kubernetes, Terraform, CI/CD pipelines

ARCHITECTURE

Medallion architecture, data lakehouse, microservices, distributed systems, REST, GraphQL

DATA & ANALYTICS

SQL, Apache Spark, Delta Lake, Microsoft Fabric, Apache Airflow, Prefect, Parquet, Avro, Elasticsearch, ClickHouse, MySQL, MongoDB, Redis, Google Spanner

OBSERVABILITY

Grafana, Loki, Prometheus, ELK Stack, GCP Cloud Logging, AWS CloudWatch

EXPERIENCE

Pyyme Digital Sweden Stockholm

Staff Software Engineer (Associate Manager) Feb '24 – present

Staff-level consultant at Pyyme, a software engineering consultancy with offices in the US, Brazil, and Sweden, leading client engagements from discovery and solution architecture through to hands-on implementation and delivery.

PROJECTS

Consulting Software Engineer, Everi Labs Jul '26 – present

Everi Labs, a Stockholm-based startup, provides an AI documentation engine for orthotics and prosthetics (O&P) clinics. Engaged as a consulting software engineer to improve code health, reduce tech-debt, and strengthen software architecture and design.

- Raised test coverage from 13% to 20% in the first four weeks, with further improvements ongoing.
- Delivered an observability stack implementation plan to improve monitoring and system visibility.
- Introduced processes to support more sustainable delivery for the team.

Lead Data Engineer, NewSpring Capital Nov '25 – Jun '26

NewSpring Capital, a US-based private equity firm, partnered with Pyyme to modernise their fragmented, siloed data spread across 15+ systems, including SharePoint, Salesforce, and other cloud services, by establishing a single, reliable source of truth across 25 years and 6TB+ of historical data, unlocking advanced analytics and AI capabilities.

- Led technical discovery, conducting interviews and workshops across Capital Formation, Finance, Holdings, and other business units to map data sources, surface operational pain points, and identify automation opportunities.

- Architected the full solution: a medallion-based data platform on Microsoft Fabric, enterprise data models for key domains (deals, portfolio companies, investors, funds), a phased 8–12 month delivery roadmap including Phase 2 predictive analytics (fund IRR forecasting, portfolio growth modelling), and a licensing and cost model, which formed the basis for project sign-off.
- Led Phase 1 implementation in a dual technical and team lead capacity, and built Bronze-zone ingestion pipelines (Spark, Airflow, Python, Delta Lake) while coordinating Silver-zone transformation and designing governance and access controls (Microsoft Fabric Workspaces, Entra, Microsoft Presidio) for audit-ready, investor-grade data controls.

Tech Lead, Småspararguiden Apr '25 – Oct '25

Småspararguiden is a Swedish consumer finance platform offering independent financial guidance. Brought in as tech lead to design and build Småsparaboten, an AI-powered financial assistant grounded in the platform's proprietary knowledge base, working directly with a co-founder to shape both product direction and technical architecture.

- Designed the end-to-end system architecture for the assistant, selecting the core stack: Llama3 as the LLM for its open-weight flexibility and local deployment capability, and ChromaDB as the vector store for semantic retrieval over the knowledge base.
- Implemented a RAG pipeline using LlamaIndex and LangChain, optimising document chunking, embedding strategies, and retrieval ranking to improve response accuracy and groundedness.
- Delivered a working alpha establishing the technical foundation for the product, enabling the team to iterate toward a production-ready assistant.
- Following the alpha, migrated the LLM from Llama3 to Gemini API, trading local deployment flexibility for improved response quality and reduced infrastructure overhead, which enabled the team to focus on product iteration rather than model hosting.

Senior Data Engineer, Dema Apr '24 – Mar '25

Dema is a commerce intelligence company that helps retailers improve sales, operations, and profitability through its proprietary AI solutions. Contracted as a senior data engineer to redesign the data ingestion pipeline that feeds these solutions.

- Proposed and implemented a new pipeline architecture, transitioning from an ad-hoc synchronous model to a modular, configurable asynchronous ETL service written in Python and orchestrated with Prefect, which improved data reliability and reduced customer onboarding time by up to 50%.
- Built a JSONPath-based transformation engine enabling custom data mapping schemas for any merchant without additional development effort, and introduced a reusable connector framework that simplified integration of new data sources.
- Implemented a data lake on S3 storing raw and transformed data in Parquet format, enabling efficient querying and supporting downstream analytics and machine learning workflows.

Google Hyderabad

Senior Software Engineer Sep '21 – May '23

Tech lead for an 8-person Capacity Planning team within Google Cloud Storage, owning team OKRs and driving technical delivery for forecasting pipelines that processed billions of records per run across GCS's global infrastructure, while leading initiatives to improve system reliability and reduce operational toil.

- Designed and implemented resilience and fault-tolerance mechanisms for the resource forecasting workflow, reducing failure rates by 90% and ensuring uninterrupted daily output for downstream processing.
- Extended the forecasting workflow with a second processing phase incorporating additional data validation, producing more accurate results and significantly reducing manual effort required to process outputs.
- Developed Python applications for the capacity forecasting workflow that generated resource requirement projections for Google Cloud Storage infrastructure at scale.

- Led a complex ownership transfer of the GCS resource planning codebase and all related software assets, coordinating migrations, testing strategies, and rollback plans across engineers in multiple teams and timezones.
- Organised and led the team's participation in Google Cloud's "Fixit" weeks, proactively identifying and logging tech-debt items that were subsequently cleared, and was recognised with peer bonuses for the effort.

EARLIER CAREER

Freelance Consultant, Lumos Labs Jul '23 – Jan '24

Designed and implemented the Lumos Auth module using OAuth, enabling seamless SSO across the Academy and Hackathons platforms by decoupling authentication from business logic. Built the Hackathons platform as an OAuth resource server using Spring Boot and Spring Data REST, and deployed the full stack on AWS with a CI/CD pipeline via GitHub, CodePipeline, and Elastic Beanstalk.

Senior Software Engineer → Software Architect, Turvo (Lineage Logistics) Nov '17 – Sep '21

Led engineering for a new generation of distributed microservices features. As Software Architect, reduced API latency by 60% and Redis memory usage by 50% through a redesign of the configurations framework. Developed an in-house caching framework which was widely adopted across the platform. Led efforts to reduce tech-debt in the company's codebase, developed an automated code-quality check tool using SonarQube.

VP of Engineering, Lookup → Software Architect, NowFloats Jul '15 – Oct '17

Led re-architecture of the Lookup platform from a monolithic stack to a distributed, AWS-based architecture, scaling it to 100K+ concurrent users and eliminating downtime. Post-acquisition by NowFloats, designed a microservices-based architecture for *ana.chat*, an enterprise conversation platform, deployable as single-tenant or multi-tenant cloud.

Lead Engineer, ZipDial (Twitter) Sep '13 – Jul '15

Took over ownership of voice-based application development and built internal tooling for ZipDial's cloud telephony platform. Created a reusable framework for telephony apps that halved development time, and worked closely with client-facing teams to translate that into 50% faster solution delivery.

Founding Engineer & Software Architect, Innoz Nov '09 – Aug '13

Joined as founding engineer while in college and built the award-winning SMS Gyan service (an SMS-based search engine) from the ground up, scaling it to serve 100M+ users and 1M+ daily queries following a partnership with India's largest telecom. Evolved the architecture from a PHP/MySQL monolith to a distributed, service-oriented system on Linode. Built and led the engineering team across product, R&D, and platform functions.