

PROFILE

French-British engineer, I **design and industrialise AI platforms built for production**. Graduate of the ISIS engineering school (INSA group, top of my class) and Google Cloud certified Professional Data Engineer, I have built end-to-end MLOps chains in demanding settings (manufacturing, healthcare, finance, cybersecurity, environment), across France and internationally.

I currently **own an AI platform for supply chain operations**: I hold its architecture and every technical decision runs through me. My work is turning data science approaches into reliable, secure, operable services, then designing the **agentic layers (ADK, MCP servers)** that make those capabilities directly usable by business teams.

Co-inventor on **two patents** covering adaptive configuration of data processing chains, and winner of **Atos' 2023 Inventor Trophy**. A former high-level athlete, I kept the discipline and the stamina from it; my dual French-British background makes me comfortable in international environments.

SKILLS

LANGUAGES & DEVELOPMENT

Python 6

FastAPI 3

REST APIs 4

SQL 4

PySpark 2

Bash / Linux 5

Git / GitHub / GitLab 6

YAML 4

Django

Java / Spring

PLATFORM, INFRASTRUCTURE & SECURITY

Docker 4

Kubernetes 3

RustFS / S3 object storage 1

CI/CD (GitLab CI, Azure DevOps, GitOps) 4

Terraform / IaC 2

Argo CD / Flux CD 2

Microsoft Azure 3

Google Cloud Platform 3

AWS (ECS, Lambda, CodePipeline) 2

Azure Kubernetes Service 2

Azure Container Apps 1

On-premise deployment 3

Application security & DevSecOps 3

Observability (Prometheus, Grafana) 3

AI, MACHINE LEARNING & MLOPS

MLOps 4

Model industrialisation 4

Classical ML / time series 4

MLflow 3

Prefect 1

Apache Airflow 3

Azure Databricks 2

Agentic systems (ADK) 1

Model Context Protocol (MCP) / FastMCP 1

LLMs / GenAI 3

RAG & vector databases 3

Azure OpenAI / AI Search / AI Speech 1

Reinforcement learning 2

Anomaly detection 3

Operations research (OR-Tools, CPLEX) 3

DOMAINS & BUSINESS KNOWLEDGE

Supply chain & production planning 1

ERP / Sage X3 1

Sales & demand forecasting 1

Production scheduling 2

Replenishment optimisation 1

Health information systems 4

Cybersecurity 3

IoT & sensor data 2

ARCHITECTURE & WAYS OF WORKING

Solution architecture 3

Cloud-native architecture 3

Platform as a service 4

Multi-tenancy & data isolation 2

Agile / Kanban 6

Technical scoping & pre-sales 2

Technical leadership & code review 2

Proof of concept → production 3

Collaborative problem solving 6

The figure indicates years of hands-on practice.

Supply chain AI platform CURRENT

Groupe Flow Line | Platform owner

Since 01/2026

Lyon, France

CONTEXT

Flow Line, historically an ERP integrator now repositioning as an **IT services company**, has launched an AI platform that **plugs into its clients' ERP** (reference connector: **Sage X3**, architecture designed to be ERP-agnostic) and **extracts their supply chain data**.

On that data foundation, the platform delivers several decision services: **sales forecasting, production scheduling and planning, anomaly detection and client prospecting**, and **replenishment optimisation**, up to raising purchase order proposals.

The core relies on **classical machine learning**, explainable, frugal, and retrainable per client. Every capability is nonetheless designed to be **driven by AI agents**, through **ADK** and **MCP servers**.

It is the group's flagship product and its main commercial growth driver.

MY ROLE

I own the platform: I hold its architecture, and every technical and product decision goes through me. The AI team is four people; use cases and data science approaches are explored upstream, and **I industrialise them**: turning a validated approach into a reliable, secure, production-grade service.

- **Architecture and technical leadership** of the platform:
 - defining the target architecture and arbitrating every technical choice,
 - a modular core per domain (forecasting, scheduling, anomalies, replenishment) on top of a shared ERP ingestion and normalisation layer,
 - **multi-tenant** design with strict isolation of each client's data.
- **Industrialisation of data science work**, from prototype to production service:
 - workflow orchestration with **Prefect** (ingestion, retraining, scheduled computation),
 - model versioning and performance tracking with **MLflow**, artifacts and datasets stored on **RustFS**,
 - packaging, testing and continuous integration,
 - monitoring, logging, alerting and failure recovery.
- **Design of the agentic layer (ADK + MCP servers built with FastMCP)**:
 - exposing each business capability as a documented, agent-callable tool (recompute a forecast, run a scheduling pass, generate a replenishment proposal...),
 - guardrails and **systematic human validation** before any side-effecting action,
 - full traceability of the agent's actions.
- **Security and compliance**: access control, secret management, isolation of client environments, auditability of processing.
- **Interface with the data science team**: upstream technical scoping of new use cases, interface contracts, code review and quality uplift.
- **Pre-sales support**: demonstrations, answers to prospects' technical and security requirements, integration effort estimates.

Artificial intelligence (AI)

MLOps

AI agents

ADK

Model Context Protocol (MCP)

FastMCP

Machine learning

Supply chain

Demand forecasting

Production scheduling

Anomaly detection

ERP

Sage X3

Python

FastAPI

Prefect

MLflow

RustFS

Docker

CI/CD

Solution architecture

Multi-tenancy

Data security

Platform as a service

Production order scheduling (agentic layer)

Groupe Flow Line | Client: APERAM

06/2026

France

CONTEXT

APERAM commissioned Flow Line to build a **bespoke planning and optimisation tool for production orders**.

The challenge: give planners a way to build, compare and adjust optimised production plans without going through a hard-to-use expert tool.

MY ROLE

The optimisation engine was built by the data science team. I joined the project to **make it agentic**: to turn it into a tool driven in natural language, fully traceable and safe.

- Designed and implemented an **agentic layer (ADK)** exposed through a **chatbot**, letting planners talk to the solver in natural language.
- Built the **tooling made available to the agent**:
 - moving and rescheduling production orders,
 - running the optimisation engine (**OR-Tools**) and comparing scenarios,
 - producing a readable summary of what each iteration changed.
- Implemented **end-to-end traceability**: logging of every agent action, change history and reversibility of decisions.
- Defined the **guardrails**: business constraint checks and human validation before a plan is applied.
- Integrated the solution into the existing environment and supported end users.

Artificial intelligence (AI)

AI agents

ADK

Chatbot

LLMOps

Operations research

OR-Tools

Scheduling

Optimisation

Python

Traceability

Manufacturing

RAG chatbot

07/2025 – 01/2026

Inetum | Client: Université Jean Moulin Lyon 3

Lyon, France

CONTEXT

Université Jean Moulin Lyon 3 wanted to improve access to information and the user experience on its **Moodle** platform, where students and teaching staff regularly need help navigating courses, finding documents or obtaining administrative information.

The assignment was to deploy a **RAG chatbot** answering in real time from the university's internal document base.

MY ROLE

- Designed and deployed an **on-premise RAG chatbot** on the university's own infrastructure, guaranteeing the security and confidentiality of internal data.
- Integrated the chatbot with **Moodle**, enabling:
 - automated lookup of teaching and administrative documents,
 - guided navigation through courses and modules,
 - general information about the university and how its programmes run.
- Worked with the IT and teaching teams to structure the document base and streamline the question-answer flow.
- Ran tests and iterations to improve answer relevance and natural language understanding.
- Set up performance tracking and monitoring to guarantee the tool's reliability and security.

Artificial intelligence (AI)

Retrieval-Augmented Generation (RAG)

Vector databases

LLMOps

GenAI

Python

FastAPI

Docker

Kubernetes

Linux

Bash

Git

DevOps

MLOps

Architecture

Knowledge Hub

02/2025 – 01/2026

Inetum

Lyon, France

CONTEXT

KnowledgeHub is an internal platform for **centralising and intelligently orchestrating knowledge**, unifying access to documents, databases and enterprise resources.

It embeds a **conversational AI agent** able to answer complex natural language queries through modules such as **RAG**, **text-to-SQL** and **SQL-to-visualisation**, making information accessible and usable without technical skills.

MY ROLE

I designed, developed and deployed **the entire platform**, except the initial API development, which I **migrated from Flask to FastAPI** for better performance and a more modern integration.

- Built an **MCP (Model Context Protocol) server** to dynamically supply contextual tools to the AI agent based on user queries.
- Deployed the platform on **Azure Kubernetes Service (AKS)** with **Terraform**, with **CD through Flux CD** (GitOps) for reliable, automated releases.
- Set up **CI pipelines in Azure DevOps**: build, test, packaging, validation.
- Orchestrated the AI components with **Kubernetes** and integrated the modules (RAG, text-to-SQL, SQL-to-viz) into a scalable architecture.
- Developed the user interface connected to the AI agent, integrated internal data sources, handled conversational context and access rights.
- Owned **infrastructure maintenance**, monitoring (metrics, logs, alerting) and the platform's continuous evolution against business needs.

Artificial intelligence (AI)

AI agents

Model Context Protocol (MCP)

Retrieval-Augmented Generation (RAG)

LLMOps

GenAI

Azure Kubernetes Service (AKS)

Terraform

GitOps

Azure DevOps

Microsoft Azure

FastAPI

Django

Python

Linux

Bash

YAML

DevOps

Solution architecture

Product innovation

Knowledge transfer application

10/2025 – 12/2025

Inetum | Client: SNCF

France

CONTEXT

SNCF wanted to **capture and pass on the know-how of its experienced staff** to new joiners, in a context of team renewal where a large share of operational knowledge is undocumented.

A **proof of concept** had shown the full chain was feasible: recording interviews with experienced staff, **automatic transcription** through **Azure AI Speech**, **summarisation** through **Azure OpenAI**, then delivery through a **RAG chatbot** backed by **Azure AI Search** and queried in natural language.

The prototype proved the idea but was not usable as it stood: no tests, no reproducible deployment, no monitoring, and uncontrolled cognitive-service costs.

MY ROLE

I came in to **industrialise the solution**, running a **complete refactor of the proof of concept** into a deployable, maintainable, production-grade application.

- **Full architectural rework**: from a monolithic prototype to a modular application where each stage of the chain (capture, transcription, summarisation, indexing, retrieval) became an isolated, testable, replaceable component.
- **Hardening of the Azure processing chain**:
 - ingestion and transcription of interviews with **Azure AI Speech**,
 - automatic generation of usable summaries through **Azure OpenAI**,
 - indexing and semantic search in **Azure AI Search**, feeding the **RAG chatbot**.
- **Production deployment on Azure Container Apps**: each component containerised, CI/CD pipelines, secret and managed-identity handling, separated environments, monitoring and failure recovery.
- **Cost control**, critical as the volume of interviews grew: rationalised cognitive-service calls, caching and batch processing.
- **Security and confidentiality**: the interviews contain sensitive internal material, hence strict access partitioning and auditable processing.
- **Quality and handover**: test coverage, code structuring and documentation so SNCF's own teams could take ownership.

IRC01/2025 – 06/2025

Inetum | Client: Crédit Agricole Technologies et ServicesLyon, France

CONTEXT

Crédit Agricole Technologies et Services engaged Inetum to speed up the analysis and reporting of internal survey results (performance, client satisfaction, business feedback) through a platform able to **centralise, analyse and automatically render data** as **PowerPoint reports enriched with charts**.

The goal was to simplify analysts' work by combining **RAG, automated visualisation** and **intelligent report generation** modules.

MY ROLE

I played a key role in **infrastructure optimisation, deployment automation** and **technical delivery to the client**.

- **AWS infrastructure optimisation:**
 - reorganised the existing **ECS clusters** (clean-up, rationalisation, better service allocation),
 - **targeted migration to AWS Lambda** for selected functions, cutting cost and complexity.
- **Deployment automation: CI/CD pipelines with AWS CodePipeline**, delivering fast, reliable, reproducible releases with no manual steps.
- **Client support & delivery:**
 - technical handover of the final codebase with documentation,
 - supported Crédit Agricole in integrating the solution into their internal architecture (environments, security, interconnections).
- **Overall productivity gains:** automation of analysis and reporting workflows to remove time-consuming manual tasks.

Sovereign AI

Atos

08/2024 – 01/2025

Montpellier, France

CONTEXT

Run in partnership with Dell, this project delivered a complete **on-premise** offering for putting AI models into production: optimised hardware infrastructure, a tailored MLOps stack, plus maintenance and governance services. It answers the critical needs of **data protection**, **low latency** and **full control** over AI models in local environments, for regulated or sensitive sectors.

MY ROLE

I contributed to the design of the Sovereign AI platform, in particular its **MLOps** layer, from architecture choices through to technical implementation.

- Built the **full MLOps infrastructure**, orchestrated with **Kubernetes**, supporting data ingestion, training, deployment and monitoring of AI models.
- Developed the platform's **core components**: ML pipeline management (Python, **Apache Airflow**), continuous model deployment (**Argo CD**), versioning and performance tracking (**MLflow**).
- Selected and integrated tooling suited to **on-premise** use, with particular attention to **scalability**, **data security** and **maintainability**.

Artificial intelligence (AI) MLOps Kubernetes Apache Airflow Argo CD MLflow CI/CD GitOps Docker Python
PySpark Data lake Infrastructure as code (IaC) Systems deployment Performance evaluation Linux Bash Git YAML

AI 4 Code

Atos

08/2023 – 08/2024

Montpellier, France

CONTEXT

AI4Code was launched to help DevOps teams without cybersecurity expertise **automate the detection and remediation of code vulnerabilities**.

The solution combines **SAST/DAST tooling (KICS)** with an **on-premise LLM** using **RAG**, able to automatically raise **corrective pull requests**.

The aim: **smoother deployments** with **stronger code security**, through a platform that is **agnostic to projects, LLMs and Git repositories** and drops into any CI/CD pipeline.

MY ROLE

I contributed to the design and development of the platform's technical infrastructure as a whole.

- Built an **on-premise infrastructure** with **Docker**, **Kubernetes**, **Airflow** and **Argo CD** to orchestrate the analysis and fix-generation workflows.
- Integrated **Ray Serve** for distributed calls to the on-premise LLM, with **RAG** generating contextual fixes from the **ChromaDB** vector store.
- Implemented the **reinforcement learning loop**: prompts are automatically refined based on whether pull requests are accepted, improving LLM performance over time.
- Handled data (code and metadata) with **MinIO** for artifact storage, **Prometheus** and **Grafana** for monitoring.
- Designed the platform's Infrastructure as Code to be **fully configurable through YAML**, allowing smooth deployment in any environment and against any Git repository.
- A cross-cutting role at the intersection of **MLOps**, **security** and **platform engineering**, strongly oriented towards **secure automation of DevSecOps workflows**.

Artificial intelligence (AI) MLOps DevSecOps Cybersecurity GenAI LLMOps Retrieval-Augmented Generation (RAG)
Reinforcement learning Ray Serve ChromaDB MinIO Prometheus Grafana Kubernetes Docker Apache Airflow Argo CD
Python YAML Linux Infrastructure as code (IaC) Lean management

Wastewater Intelligent Network (WWIN)

Atos | Client: Scottish Water

12/2023 – 05/2024
Scotland, United Kingdom

CONTEXT

Scottish Water, Scotland's main drinking water provider, launched a transformation programme to bring intelligent management to its infrastructure.

The goal was to move from reactive to proactive, predictive management of the wastewater network using real-time data, **reducing pollution incidents and flooding**, limiting environmental impact and meeting SEPA's regulatory requirements.

MY ROLE

WWIN delivered an intelligent wastewater management network: progressive IoT sensor rollout, real-time data collection and exploitation, integration into a cloud platform with visualisation, alerting and predictive models. I owned the **MLOps** side.

- Automated the **machine learning pipelines**: data ingestion, processing (**PySpark**), model training/deployment and monitoring on **Azure Databricks**, with **MLflow** for experiment tracking.
- Set up **CI/CD with Azure DevOps** to ensure fast, reliable delivery of developments.
- Worked within a multidisciplinary team: hydrologists, data scientists, data engineers, a data architect, and myself as MLOps engineer.

Artificial intelligence (AI) MLOps MLflow Azure Databricks Microsoft Azure Azure DevOps PySpark Python SQL
DevOps IoT Data engineering Performance management Agile English-speaking environment

FIDAMIA

Atos | Client: Aix-Marseille-Provence Metropolitan Authority

04/2023 – 06/2023
Marseille, France

CONTEXT

The Aix-Marseille-Provence metropolitan authority commissioned Atos to deliver an AI solution improving the handling of citizen reports submitted through the "Ma Métropole Dans Ma Poche" mobile app.

The aim was to **automate the categorisation of reported incidents** (cleanliness, roads, signage, etc.) through contextual image analysis, **shortening response times** for technical services and **improving the citizen journey**.

MY ROLE

Atos deployed a lightweight but robust infrastructure for industrialised AI model deployment on two virtual machines, with a full MLOps chain. I contributed to the **network and security infrastructure**.

- Implemented **reverse proxies with Traefik and Kong** for secure traffic management, request routing and controlled exposure of the AI services.
- Configured **certificates**, HTTPS redirection and authentication/access restriction mechanisms.
- Outcome: reliable, secure and performant access to the AI services in the production environment.

Artificial intelligence (AI) Network security Network architecture Traefik Kong Kubernetes MLOps

IoT dataset selection through reinforcement learning

03/2023 – 06/2023

Atos

Montpellier, France

CONTEXT

Design and development of an **MLOps** platform for **IoT data processing** geared towards anomaly detection. It is the work for which I received **Atos' 2023 Inventor Trophy**.

Its main innovation is the use of **reinforcement learning** to dynamically select training datasets according to usage context and expected performance.

The project led to a **European and US patent filing** (co-filed with a colleague), covering the adaptive data management approach for training ML models in distributed, constrained environments.

MY ROLE

- Led the **MLOps architecture design** and developed the full **data processing chain**, from IoT ingestion through to model deployment.
- Designed and integrated the **reinforcement learning** module driving dataset selection in production.
- Contributed to **drafting the patent** alongside Atos' expert team, formalising the technical innovation developed.

Artificial intelligence (AI)

Reinforcement learning

MLOps

Machine learning

Anomaly detection

IoT

R&D

Apache Airflow

Docker

Python

Kanban

MLSecOps platform

01/2023 – 03/2023

Atos

Montpellier, France

CONTEXT

An internal Atos R&D assignment to build an **MLSecOps** platform securing the full chain of data processing and AI model deployment.

The solution takes a novel approach to auditability and traceability of critical steps by recording integrity hashes on the **Ethereum blockchain** through **smart contracts**.

MY ROLE

I was responsible for developing and integrating the whole platform (smart contracts excepted).

- Built **secured MLOps chains**, from data ingestion through to model deployment.
- **Integrated with the Ethereum blockchain** to verify the integrity of critical steps (datasets, models, artifacts).
- Created an **integrity proof mechanism** through hashes published in **Ethereum smart contracts**.
- Deployed a secured on-premise Docker infrastructure to run the pipeline stages.
- Implemented automated verification workflows guaranteeing object integrity across the whole lifecycle.

Artificial intelligence (AI)

MLOps

DevSecOps

SecOps

Cybersecurity

Pipeline integrity

Data integrity

Blockchain

Ethereum

Solidity

Smart contracts

Web3

Docker

Python

CONTEXT

At Atout Majeur Concept, a software publisher for hospitals, I worked on **logistics optimisation** scenarios to improve the management of **patient portering** flows.

The objective was to model porter assignment so as to minimise patient waiting times and optimise routes within the facility.

MY ROLE

- Designed and implemented **mathematical optimisation models** based on time, availability and priority constraints.
- Used **CPLEX Studio** to formulate and solve these complex **mixed-integer linear programming (MILP)** problems, applied to real hospital organisation situations.
- Proposed concrete improvement scenarios for the **operational management of porters**.

Operations research

Mathematical modelling

Linear programming

CPLEX

Python

Health information technology

Healthcare services management

EDUCATION

Master's-level engineering degree (Diplôme d'ingénieur), ISIS Engineering School (INSA group)

2020 – 2023
Castres, France

Computer science and health information systems, AI & Big Data specialisation. **Top of the class.**

CPGE BCPST (intensive preparatory programme), Lycée Champollion

2018 – 2020
Grenoble, France

Biology, chemistry, physics, earth sciences.

Baccalauréat, science stream, Lycée Jean Moulin (elite sport programme)

2018
Albertville, France

Highest honours (mention Très Bien).

CERTIFICATIONS

Professional Data Engineer, Google Cloud

2023

Certified August 2023.
Credential ID: TiA22q

PATENTS & AWARDS

Method for configuring a data processing chain, Patent (United States)

2024

Adaptive configuration of a data processing chain: reinforcement-based selection of the training dataset according to a reward tied to the model's accuracy score. Co-inventor.
US2024412073A1, issued 12/12/2024

Procédé de paramétrage d'une chaîne de traitement de données, Patent (Europe)

2024

European counterpart of the above. Co-inventor.
EP4475044A1, issued 11/12/2024

Inventor Trophy, Atos

2023

Internal award recognising a contribution to MLOps innovation (IoT data processing platform driven by reinforcement learning).

PROJECTS & INVOLVEMENT**Mentoring, AFEV**

2021 – 2022

Mentored a secondary school pupil from a disadvantaged background: academic support, cultural exposure, guidance on study choices.

Synthetic relational data, Supervised project, in collaboration with Accenture

2022

Three-month research project on generating synthetic multi-table relational data: state of the art (GANs, CTGAN, autoencoders) and improvement avenues.

LANGUAGES**English**, Native**French**, Native

REFERENCES (3)**André Polaszek** (on the Atos engagement (08/2025))

Nelson is approachable, capable and professional. I worked with him on the development of an MLOps platform, where he took on the tech lead role brilliantly. I recommend him wholeheartedly for any project in that field.

Léa Duran (on the Scottish Water engagement (07/2025))

Nelson demonstrated his MLOps expertise on a large-scale project for a major national water utility. He operationalised the AI pipelines feeding their IoT platform (4,000 sensors). He is a pleasure to work with, thanks to his team spirit and his autonomy.

Yannick Lecroart (mentor at Atos (07/2025))

Nelson shows an unwavering determination to learn and improve, standing out in key areas such as data science, generative AI, microservices, automation and DevOps. His rare combination of technical skill, appetite for learning and innovation makes him an outstanding professional.