

YOHANS (JOHN) BEKELE

Gdansk, Pomeranian, 81-350 | +48 739 693 705 | yohansdemisie@gmail.com

[Portfolio-website](#) | [GitHub](#) / [LinkedIn](#)

PROFESSIONAL SUMMARY

AIOps Engineer with 4+ years of progressive experience across AIOps and automation, site reliability, data analytics, and full-stack development, now focused on designing and deploying agentic AI systems in production at enterprise scale for MAIA (internal gig with TR Labs), an AI-native deal intelligence platform for private equity, corporate development, legal, and advisory teams featuring multi-format document ingestion with PII redaction and M&A taxonomy classification, hierarchical chunking, pgvector RAG, and Claude-based AI workflows for document-grounded Q&A, summarization, risk extraction, company research, fact checking, early-alert analysis, Native connection to VDR for seamless ai ingestion, and automated report generation.

Concurrently, in role as Product Support Specialist, built an internal application that's helps aggregating knowledge base from different resource like wiki, salesforce, azure tickets, snowflake and got through ETL process to then used as debugging tool and it's robust REST API integrated in production application. Earlier roles in Service Reliability Engineering and data migration built a foundation in production monitoring, incident response, and automation, with 10,000+ incidents resolved and 99.8% production uptime sustained.

Skills span Python (FastAPI, Pandas, NumPy, pytest), AWS (ECS Fargate, Aurora PostgreSQL with pgvector, Bedrock, Pulumi IaC, 20+ CI/CD pipelines, zero-trust security), the GenAI stack (prompt engineering, chunking and embedding strategy, prompt and semantic caching, RAG, agentic orchestration) and AI workflow testing, and modern web development (React, Next.js, TypeScript).

TECHNICAL SKILLS

Languages: Python, TypeScript, JavaScript, SQL, Bash, YAML

AI/ML & NLP: Claude agent SDK, Bedrock agent core, Anthropic API, LangChain, RAG, embeddings, semantic search, Information Retrieval, Information Extraction, knowledge graphs, agentic AI orchestration, prompt engineering, structured extraction, document processing pipelines

Data & Scientific Python: Pandas, NumPy, SQL/ETL pipelines, Alembic

Testing: pytest, regression / integration / acceptance test suites, Locust load testing, LLM-provider failover testing, prompt-injection and PII-redaction security testing, CI/CD quality gates

Cloud (AWS): ECS Fargate, Aurora PostgreSQL, pgvector, DynamoDB, ElastiCache, Bedrock, SageMaker, S3, CloudFront, ALB, IAM, VPC, KMS, WAFv2, GuardDuty, Security Hub, Secrets Manager, CloudWatch

IaC & CI/CD: Pulumi (Python), Terraform, CloudFormation, GitHub Actions (20+ workflows), Docker, ECR, production approval gates, automated rollbacks, feature flags

Monitoring & SRE: New Relic, Datadog, CloudWatch, Prometheus, Grafana, ELK Stack, incident response, root cause analysis, SRE practices

Databases: PostgreSQL, pgvector, MySQL, MS SQL Server, MongoDB, Aurora, DynamoDB, Redis

Full Stack: FastAPI, Next.js, React, Node.js, Express, REST APIs, Tailwind CSS, Nginx

PROFESSIONAL EXPERIENCE

AIOps Engineer (Internal Gig) — MAIA, AI-Native Deal Intelligence Platform

Thomson Reuters Labs | Gdansk, Poland | December 2025 – Present

- Engineered and deployed production AI based workflows powering an AI-native deal intelligence platform: document-grounded Q&A, summarization, risk extraction, **company research**, **fact checking**, **early-alert analysis**, and **automated report generation** — using Bedrock and Anthropic API with SSE streaming, prompt engineering, generation parameter tuning (temperature, top_p, max_tokens), **Anthropic prompt caching** for cost and latency

reduction, LLM-provider failover (Bedrock ↔ Anthropic), and structured extraction from unstructured deal documents.

- Designed and deployed the production **RAG and document processing pipeline**: multi-format document ingestion, PII redaction, M&A taxonomy classification, **hierarchical chunking**, pgvector embeddings for semantic search over multi-document legal and financial corpora, and **semantic caching** to cut embedding/LLM cost on repeated/similar queries.
- Integrated **VDR connectors** for ingesting deal documents directly from client repositories and external virtual data rooms, enabling seamless workflow from data source to indexed, AI-ready content.
- Built an **Information Extraction** layer using GenAI-driven entity and relation extraction to construct a background knowledge graph from unstructured deal documents, supporting entity-centric search and document-relationship analysis.
- Built the AWS based platform from architectural blue to practice with Pulumi (Python) to host the AI workloads: 7 microservices (5 FastAPI + 2 Next.js) on ECS Fargate across 4 environments, dual Aurora PostgreSQL with pgvector for embeddings, ElastiCache Redis, CloudFront + WAF, ALB with path-based routing, S3 with KMS encryption, and full VPC networking.
- Authored 20+ GitHub Actions CI/CD workflows with production approval gates, service-aware reviewer routing, SHA-pinned action security, workflow security scanning, automated rollbacks, feature flags, and Claude-powered release notes.
- Designed automated testing for production AI workflows: regression, integration, acceptance, LLM-provider failover, prompt-injection, and PII-redaction security suites; Locust-based load testing with SLO validation; CI/CD quality gates that run after every deploy.
- Led a production account migration end-to-end: new Pulumi stacks, S3 state backend with DynamoDB locking, and authored 6 verification audits covering IAM, networking, databases, services, deployment workflows, and security posture.
- Implemented zero-trust security posture: WAFv2 in BLOCK mode, GuardDuty with Fargate runtime monitoring, Security Hub (CIS + AWS Best Practices), KMS encryption, S3 deny policies, and per-service secret isolation.
- Resolved production incidents across the AI and platform stack: secret injection failures, PostgreSQL race conditions, Nginx redirect loops, and Bedrock-to-Anthropic API migration.

Product Support Specialist

Thomson Reuters | Gdansk, Poland | August 2024 – Present

- Built an AI-powered **Knowledge Base management system** for the application support team using **RAG and agentic AI**: agents extract and synthesize information from multiple internal sources, expose it through a robust REST API designed for integration with legacy applications, and enable end users to self-diagnose issues before opening tickets.
- Resolved **10,000+ production incidents** via ServiceNow and JIRA, with structured root cause analysis, post-incident reviews, and runbook authoring; recurring incident patterns fed back into the knowledge base.
- Detected outages and client-side application faults early using **clustering-based anomaly detection** over user activity telemetry.
- 99.8% system uptime achieved through proactive monitoring, alert tuning, and preventive maintenance across enterprise applications.
- Automated daily monitoring, health checks, and operational workflows in Python, Bash, and SQL — reducing manual effort by 15+ hours per week.
- Supported database operations and migrations across MS SQL Server, PostgreSQL, and MySQL.

Freelance AI / Full Stack Developer

Upwork & freelance.com | Remote | July 2024 – December 2025

- Built multiple end-to-end AI and web applications across finance, education, legal, real estate, agriculture, and infrastructure automation domains, using React, Next.js, TypeScript, Python/FastAPI, Node.js/Express, PostgreSQL, MongoDB, Docker, and cloud deployment workflows.
- Developed LLM-powered application features including financial news analysis, entity extraction, sentiment analysis, stock-signal generation, AI tutoring workflows, prompt design, response handling, and evaluation criteria for generated outputs.
- Designed and implemented full-stack product architectures, including REST APIs, authentication flows, database schemas, background processing, file handling, user dashboards, admin panels, and responsive frontend interfaces.

- Built infrastructure automation tooling that translates visual user input into cloud infrastructure-as-code, including drag-and-drop AWS resource modeling, Pulumi/Terraform code generation, GitHub OAuth integration, and repository-oriented workflows.
- Implemented production-style engineering practices across projects, including containerization with Docker, API testing, reusable frontend components, backend service separation, environment configuration, CI/CD setup, and deployment troubleshooting.

Service Reliability Engineer

Infosys | Poland | April 2022 – June 2024

- Operated production enterprise applications on a 24/7 support model, ensuring high availability and SLA compliance through proactive monitoring with New Relic, Datadog, ELK Stack, and custom dashboards.
- Reduced MTTR by 45% through improved incident triage, structured root cause analysis, and escalation workflows in ServiceNow and JIRA.
- Wrote and maintained Python and Bash automation for routine operational work — log analysis, health checks, alert deduplication, report generation, and ticket auto-triage — improving team efficiency by 30%.
- Built and scheduled **cron jobs and workflow automations** for daily and weekly operational tasks: scheduled health checks, batch job monitoring, log rotation, backup verification, and metric aggregation.
- Configured alerting thresholds, on-call rotations, and runbook automation to reduce alert fatigue and ensure timely escalation of true incidents.
- Supported database migrations and integrity verification across MS SQL Server, PostgreSQL, and MySQL, including pre-migration schema validation and post-migration reconciliation.
- Authored runbooks and operational documentation for incident response, recurring issues, deployment workflows, and on-call handoffs.

Data Analytics / Migration Specialist (Contract)

Cisco | Krakow, Poland | August 2022 – February 2023

- Automated ETL processes using Python, Bash, and SQL — improving data processing efficiency by 30% and replacing fragile manual workflows with reusable, parameterized pipelines.
- Designed and executed data migration workflows between heterogeneous source and target systems, with stage-by-stage reconciliation, row-count verification, and rollback safeguards to guarantee record integrity.
- Developed data validation and cleansing routines — schema checks, referential integrity, deduplication, and type normalization — ensuring accuracy and compliance before data was promoted to target systems.
- Built analytical dashboards and reports providing actionable insights to business stakeholders for compliance, audit, and business continuity decisions.
- Collaborated with cross-functional teams — business analysts, source-system owners, and target-system engineers — to coordinate cutover windows and minimize business disruption during migration.

CERTIFICATIONS

AWS Certified Cloud Practitioner

EDUCATION

Bachelor of Science in Computer Science

WSB University | Torun, Poland | 2026

LANGUAGES

English (Fluent)