

Maksim Z.

Architect/Team-lead CTO / Data & AI Architect / Lead Data Engineer

SUMMARY

- Over 15 years of expertise in data architecture, ETL pipelines, and AI/ML-driven analytics, with a PhD in Mathematics emphasizing analytical rigor.
- Proficient in Python, SQL/PLSQL, Java, Golang, and modern frameworks like FastAPI, LangChain, and LangGraph, delivering scalable AI and data engineering solutions.
- Led cloud-native microservices and event-driven architectures using Docker, NATS, and CI/CD pipelines across AWS, GCP, and IBM Cloud

TECHNICAL SKILLS

Main Technical Skills	Python (3 yr.), FastAPI (3 yr.), OpenAI (3 yr.), LangChain (3 yr.), PostgreSQL (3 yr.)
AI Tools & Assistants	DeepSeek, Gemini (1 yr.)
Programming Languages	Go, Python (3 yr.), R
Python Libraries and Tools	aiohttp, Pydantic, pytest
AI & Machine Learning	AWS SageMaker, Faiss, LangChain (3 yr.), LangGraph (1 yr.), LLaMA, Llamaindex, ollama, OpenAI (3 yr.), PandasAI, Vertex AI, Whisper
Python Frameworks	FastAPI (3 yr.), Prefect
.NET Platform	.NET
JavaScript Frameworks	Node.js
Data Analysis and Visualization Technologies	AWS Athena (1 yr.), Azure Data Factory (1 yr.), Data Scraping (1 yr.), Domo, PandasAI, Power BI (2 yr.), Presto S3 (1 yr.), Tableau
Databases & Management Systems / ORM	AWS Redshift (1 yr.), ChromaDB, IBM DB2 (1 yr.), Microsoft SQL Server (3 yr.), MongoDB, MySQL, Netezza (1 yr.), Oracle Database (7 yr.), PL/SQL (3 yr.), PostgreSQL (3 yr.), Redis, Snowflake, SQL (3 yr.), SQLAlchemy, SSIS (2 yr.)
Cloud Platforms, Services & Computing	Firestore Cloud Messaging (1 yr.)
Amazon Web Services	AWS Athena (1 yr.), AWS Bedrock, AWS CloudTrail, AWS Redshift (1 yr.), AWS SageMaker
Azure Cloud Services	Azure Arm templates (1 yr.), Azure Data Factory (1 yr.)
Google Cloud Platform	Google BigQuery

Third Party Tools / IDEs / SDK / Services	Adobe LiveCycle
Methodologies, Paradigms and Patterns	Agile, microservices (1 yr.)
Deployment, CI/CD & Administration	Azure Arm templates (1 yr.), GitLab CI
Version Control	BitBucket, Github Actions
Virtualization, Containers and Orchestration	Docker Compose
SDK / API and Integrations	FastAPI (3 yr.)
QA, Test Automation, Security	Mock, Postman, pytest, Unit Testing
Mail / Network Protocols / Data transfer	NAT (1 yr.), WebSockets (2 yr.)
UI/UX/Wireframing	Prototyping (1 yr.)
File Systems, Storage	XML (1 yr.)
Other Technical Skills	AutoGen, Python Requests Library

WORK EXPERIENCE

Senior Backend Developer | Tech Lead (AI Voice-to-Voice Telephony Verification System)

Duration: Oct 2025 – Jan 2026

Summary:

- Designed and implemented an AI-driven voice-to-voice telephony system for automated client identity and account status verification
- The system autonomously initiates outbound calls, conducts verification dialogs, and captures structured verification results, supporting real-world telephony scenarios including voicemail detection and IVR navigation
- It uses a Retrieval-Augmented Generation (RAG) layer to fetch relevant client and contextual data during live calls, enabling dynamic prompt adaptation in real time with full end-to-end automation including monitoring and escalation flows

Responsibilities:

- Designed end-to-end architecture for automated voice verification workflows;
- Implemented outbound voice call handling with IVR and voicemail detection;
- Built real-time dialog management with dynamic prompt adaptation;
- Integrated RAG pipelines to retrieve client and account context during calls;
- Implemented structured result extraction and verification outcome storage;
- Designed monitoring, confidence scoring, and escalation mechanisms for uncertain cases;
- Implemented voice interfaces for conversational interaction.



Technologies: Python, FastAPI, OpenAI Realtime, RAG, Vector Databases, Telephony APIs, Docker

Senior Backend Developer | Tech Lead (AI-Based Buying Committee Identification Platform)

Duration: Oct 2025 – Jan 2026

Summary:

- Designed and implemented an AI-powered system for identifying a company's Buying Committee within an existing contact base
- The platform builds structured company profiles by combining PDF parsing, AI-based content analysis, and data aggregation from multiple external sources via APIs and scraping
- It uses a custom LLM-driven algorithm with prompt chains to infer roles, influence levels, and relationships within the Buying Committee, followed by automated validation and provides analytical insights through interactive dashboards

Responsibilities:

- Designed system architecture for Buying Committee identification workflows;
- Implemented PDF parsing and AI-based document analysis pipelines;
- Built data aggregation and scraping mechanisms via APIs and external sources;
- Developed a custom LLM-based prompt-chain algorithm for role and influence detection;
- Implemented validation logic to ensure result accuracy and consistency;
- Designed analytics layers with interactive charts and visualizations for end users.

Technologies: Python, FastAPI, LLMs, LangChain, PDF Parsing Tools, Web Scraping, APIs, Vector Databases, BI & Visualization Tools

Senior AI Developer | Tech Lead (AI Platform for Media Analytics)

Duration: Oct 2024 – Oct 2025

Summary:

- Developed and maintained an AI-driven platform component for orchestrating media data workflows using a tool-based LLM agent
- The system enables direct ingestion of media content from external sources including YouTube channels, search results, and individual content URLs, interactive exploration of project datasets through table-based querying, and generation of analytical reports with optional LLM-powered processing
- The agent pipeline was designed for deterministic behavior, schema correctness, and reliable tool orchestration across complex workflows

Responsibilities:

- Designed, built, and owned the LLM agent pipeline end-to-end;
- Implemented tool-based agent orchestration for media data ingestion and analysis workflows;
- Integrated external content sources, including YouTube channels, search APIs, and direct content URLs;
- Enabled interactive exploration of project tables using structured, ag-grid-style query patterns;
- Developed analytical report generation with optional LLM-based post-processing;
- Stabilized and standardized Jinja-based prompt templates to ensure deterministic tool selection and schema-valid outputs;



- Maintained an evaluation and regression testing harness to validate agent behavior, expected tool execution paths, and output correctness.

Technologies: Python, LangChain, LangGraph, Langfuse, OpenAI, OpenRouter, Gemini, Jinja Templates, APIs

CTO | Data AI Architect (Hell & Heaven Travel – AI-Powered Personal Travel Concierge)

Duration: Oct 2024 – Oct 2025

Summary:

- Designed and implemented an AI-powered personal travel concierge that helps users plan trips, receive personalized recommendations, and interact via chat and voice
- The system is built on a multi-agent architecture where specialized AI agents collaborate to handle itinerary planning, recommendations, and user communication across multiple channels

Responsibilities:

- Designed a multi-agent AI system using integrated LLMs;
- Architected cloud-native backend and application infrastructure;
- Developed a web application with messaging-platform chatbot integration;
- Implemented voice interfaces for conversational interaction.

Technologies: Python, FastAPI, OpenAI, LangChain, Docker, AWS, WebSockets, Messaging APIs

CTO | Data AI Architect (OnlyAsk – AI-Powered Teacher Assistant)

Duration: Mar 2025 – Oct 2025

Summary:

- Built an AI-powered assistant for teachers that supports lesson preparation, student interaction, and contextual Q&A
- The system combines multi-agent LLM orchestration with real-time voice and chat interfaces and a Retrieval-Augmented Generation pipeline to ensure accurate, context-aware responses

Responsibilities:

- Designed a multi-agent AI architecture using integrated LLMs;
- Developed a cloud-native web application and chatbot;
- Implemented real-time voice streaming interfaces;
- Built a Retrieval-Augmented Generation (RAG) pipeline to improve contextual accuracy.

Technologies: Python, FastAPI, OpenAI, LangGraph, LangChain, PostgreSQL, Docker, WebSockets

Data Engineer | Senior Backend Engineer (NDA Project – Enterprise Legacy Backend & CI/CD Modernization)

Duration: Sept 2024 – Oct 2025

Summary:

- Supported and modernized large-scale enterprise backend systems built on legacy platforms



- The project focused on stabilizing core business logic, improving deployment reliability, and introducing automated CI/CD pipelines

Responsibilities:

- Supported large-scale backend systems using PL/SQL and Java;
- Worked with legacy platforms including Dragon Designer;
- Designed and implemented CI/CD pipelines using GitLab CI/CD;
- Integrated AWS services to ensure reliable automated deployments.

Technologies: PL/SQL, Java, GitLab CI/CD, AWS, Oracle, Docker

Senior Data Engineer (NDA Project – Data Warehouse & BI Platform)

Duration: Nov 2024 – Sept 2025

Summary:

- Designed and implemented a corporate Data Warehouse and analytics platform to support business reporting and decision-making
- The solution included ETL pipelines, relational storage, and BI dashboards

Responsibilities:

- Designed and built DWH systems on PostgreSQL and MS SQL;
- Developed ETL workflows using SSIS and Azure Data Factory;
- Implemented analytical dashboards using Power BI.

Technologies: PostgreSQL, MS SQL, SSIS, Azure Data Factory, Power BI

Data Engineer (NDA Project – XML to DWH ETL System)

Duration: Feb 2025 – Aug 2025

Summary: Built a dedicated ETL system for transferring and transforming structured XML data into a relational Data Warehouse optimized for reporting and analytics.

Responsibilities:

- Built Data Warehouse solutions on MS SQL;
- Developed ETL pipelines for XML-to-database data ingestion;
- Ensured data validation and consistency across transformations.

Technologies: MS SQL, SSIS, XML, SQL

Senior Data Engineer (NDA Project – IBM Cloud DWH Recovery & Refactoring)

Duration: Sept 2024 – Oct 2024

Summary: Restored and refactored an existing Data Warehouse hosted on IBM Cloud to improve performance, stability, and maintainability after operational degradation.

Responsibilities:

- Restored broken Data Warehouse pipelines;
- Refactored schemas and ETL logic for better maintainability;
- Optimized performance in IBM Cloud environment.

Technologies: IBM Cloud, DB2, SQL, ETL Tools



Senior Data Engineer (NDA Project – AWS Data Ingestion & Reporting Platform)

Duration: Jan 2025 – Oct 2025

Summary: Engineered a scalable data ingestion and reporting platform on AWS to support analytics and business intelligence workloads.

Responsibilities:

- Engineered data ingestion pipelines using AWS services;
- Built workflows using Redshift, S3, and Athena;
- Designed optimized data models and stored procedures;
- Ensured consistent and reliable reporting outputs.

Technologies: AWS Redshift, S3, Athena, SQL, Python

VP Engineering – Core Technology - Act Trader Technologies (Act Trader Technologies – Core Trading Platform Modernization)

Duration: Mar 2024 – Aug 2024

Summary: Led modernization of a legacy trading platform by migrating to a cloud-native, event-driven microservices architecture with improved scalability, observability, and fault tolerance.

Responsibilities:

- Introduced cloud-native microservices architecture;
- Refactored 12+ core trading services;
- Implemented NATS-based event-driven communication;
- Improved fault tolerance through containerization and CI/CD;
- Led distributed engineering teams across multiple time zones.

Technologies: Docker, NATS, CI/CD, Cloud Platforms, Git

CTO | Tech Lead | Senior DWH Developer - DataMola (DataMola – Enterprise Data & AI Solutions)

Duration: Nov 2015 – Feb 2024

Summary: Led and delivered enterprise-grade data engineering and analytics solutions for clients across finance, logistics, telecom, and e-commerce industries.

Responsibilities:

- Led 20+ full-cycle data engineering projects;
- Designed enterprise data lakes and analytics platforms;
- Built event-driven and self-healing ETL pipelines;
- Integrated ML models for forecasting, fraud detection, and behavior analysis;
- Established coding standards, CI/CD processes, and observability practices;
- Reduced ETL maintenance costs by 40% through reusable components.

Technologies: PostgreSQL, Oracle, Python, ML Frameworks, CI/CD, Cloud Platforms

Team Lead – Risk & Reporting - SoftClub (SoftClub – Risk & Regulatory Reporting Platform)

Duration: Jan 2009 – Nov 2016

Summary: Developed a regulatory-compliant data warehouse for financial and governmental



reporting with full auditability and automated reporting pipelines.

Responsibilities:

- Led development of regulatory-compliant DWH on Oracle / Exadata;
- Translated regulatory requirements into automated data flows;
- Managed agile delivery across development and analyst teams;
- Coordinated with finance and compliance stakeholders.

Technologies: Oracle, Exadata, PL/SQL, ETL Tools

EDUCATION

- **PhD in Mathematics (p-Adic Analysis)**
Belarusian State University
- **Master's degree in Mechanics & Mathematics**
Belarusian State University
- **Postgraduate Finance & Banking Courses**
Belarusian State Economic University

