TINASHE MAKAMURE

SOFTWARE ENGINEER

aws certified Cloud Practitioner FOUNDATIONAL · Microsoft CERTIFIED AZURE FUNDAMENTALS · aws certified Solutions Architect ASSOCIATE · aws certified Developer ASSOCIATE · CERTIFIED kubernetes ADMINISTRATOR · Microsoft CERTIFIED FABRIC DATA ENGINEER ASSOCIATE

CONTACT

- ✉ tinashe@tinashemakamure.co.za
- Johannesburg, Gauteng
- tinashemakamure.co.za
- linkedin
- github

EDUCATION

2017 - 2020

UNIVERSITY OF JOHANNESBURG

- BEng Electrical and Electronic Engineering
- BSc Information Technology

LANGUAGES

- English (Fluent)

SKILLS

- AWS Cloud
- Kubernetes
- Helm Charts
- C#, JavaScript, Python, React, React Native, SQL
- RDMS
- Microsoft Fabric
- Linux
- Machine Learning
- AI Product Solutioning
- Software and Solutions Architecture
- Product Specialist
- Client communication
- People Management

PROFILE

An ageless, effulgent, and self-deprecating existence, presently exhibiting exceptional intrinsic prowess as a software engineer in the Southern region of Africa. Devoted to life-long learning and continual self-improvement, this sophisticated existence considers <code>this.self</code> to be one of the fastest developing and supremely dynamic young talents, not just in the solar system, but in the entirety of the milky way. Fields of interest include, without limiting thereto: solutions architecture, software architecture and development, cloud computing, machine learning and AI, big data, block chain technologies, financial markets, string theory, aliens, and a pinch of politics.

WORK EXPERIENCE

Synthesis Software, Johannesburg

JAN 2021 - OCT 2022

Junior Cloud Engineer (various clients)

JAN 2021 - APR 2021

On the cloud team, we built, configured, and secured cloud architectures for various esteemed clients in the financial services and retail sectors. I made use of Terraform as the IaC (Infrastructure as Code) tool of choice to carry out my tasks in building the desired architecture. We extensively made use of Terraform Cloud to automate and continuously deliver our builds to the client's AWS account. Tech Stack:

- Terraform
- AWS CLI
- AWS Console

Software Engineer (TransUnion)

MAY 2021 - JUL 2022

On the TransUnion project, I was part of a two-person team where we built some APIs for the client to enable new processes and services that they introduced to modernise their product offering. I was actively involved in the entire development process from requirements gathering, to system design, development, testing, deployment, maintenance, and support of the APIs. We split the APIs to be built between the two team members, with each member fully responsible for building a set of APIs from start to production, an enormous responsibility for each of the members.

The technology stack used was:

- C# .NET (Backend)
- PostgreSQL, Oracle SQL, Redis (Databases, Caching)
- Nginx (Reverse Proxy, Load balancing)
- Linux – RHEL8 (Deployment Environment)
- TLS certificates (Inter-Service Authentication/Authorisation, Encryption in Transit)

Software Engineer (AlumniServ)

AUG 2022 - OCT 2022

On the Aluminserv project, I was part of a five-person team, with three of the members responsible for the front-end, and myself and the tech-lead responsible for the back end. I was responsible for relational database design and implementation, writing and exposing APIs to interact with the stored data, and supporting the front-end developer experience as they used the APIs.

The technology stack used was:

- C# .NET (Backend)
- MSSQL
- Confluent Kafka
- Linux – RHEL8 (Deployment Environment)

Entelect Software, Johannesburg

NOV 2022 - JUL 2025

DevOps Engineer (Standard Bank)

FEB 2023 - OCT 2023

On this project, our team was responsible for the platform and DevOps facilitation and support for Standard Bank business banking teams. We created and maintained pipelines, AWS cloud infrastructure, managed and maintained the Kubernetes clusters where the applications were deployed across multiple environments such as Development, Test and Production. We configured and provided monitoring and observability functionalities for the applications running in the cluster and supported developer teams with deployments to production as well as troubleshooting any cloud/DevOps related issues they encountered. My duties in the team included, without limiting thereto:

- Provisioning AWS cloud infrastructure and security remediations in IAM using Infrastructure as code with Terraform and Terragrunt.
- Developed, maintained and deployed helm charts onto Kubernetes for applications that needed to be deployed on the Kubernetes cluster.
- Enhanced monitoring and observability capacity into the Kubernetes cluster, by utilising technologies like Prometheus, Thanos, Grafana, Alert manager.
- Enabled monitoring of our log processing tool, Fluentbit using Prometheus and Grafana.
- Identify potential areas of improvement or limitations in the project and create tasks on the Jira board.
- Did technical presentations on the technologies we used, their capabilities, and how we used them in the context of our project to spread this knowledge within the team.
- Did production and non-production support of several teams that used the AWS cloud platform and the Kubernetes clusters we managed.
- Supported developer teams with their deployments to production environment using AWS cloud and Argo CD.
- Mentoring junior engineers on the team.
- Participating in agile rituals and facilitating scrums, retros, PI planning session.

The technology stack used was:

- AWS Cloud
- Terraform

- Kubernetes
- Helm Charts
- Prometheus
- Thanos
- Alert Manager
- Grafana
- ArgoCD
- YAML
- Fluentbit
- Linux

Software Engineer (Osiris Trading - Betway Africa)

NOV 2023 - JUL 2025

On this project, I was part of a 6-person team to do development work in C# .NET, python, T-SQL, Azure and Microsoft Fabric, React with NextJs.

The project had multiple sub projects in it, including but not limited to:

The Sox Audit Application

During this sub-project, we developed a solution to automate the client's Sox Audit process. We utilised Azure Logic Apps to orchestrate the audit process. The logic App was powered by Azure function written in C# .Net as well as SQL stored procedures on-prem that we connected to via an On-Premises gateway.

During this project, after we had achieved a stable product, I created Infrastructure as Code for our client, using ARM (Azure Resource Manager) templates so we can easily reproduce the infrastructure in any new environment.

This later proved to be useful as other companies under the same group as our client later needed the product, and using IaC, it was quick and easy to get the product up and running in their environments.

T-SQL stored procedure optimisations

On these sub-projects, the client had a lot of their business logic in T-SQL stored procedures, and some of the procedures were becoming bottlenecks for their business processes. We were then tasked with going through the stored procedures, documenting any inefficiencies and suggesting optimisations to make the procedures more efficient in execution.

ETL processes with SSIS (On-prem)

Some of the stored procedures mentioned above were also part of a bigger ETL process that ran on-premises. Such ETLs were orchestrated by means of SSIS packages. On such projects, were tasked with extending the functionality of the ETL process to add new functionality and re-deploy the SSIS packages to the respective servers.

ETL Processes with Fabric Data pipelines (Cloud – Microsoft Fabric)

As part of their migration from On-Premises to cloud with Microsoft Fabric, on of the sub-projects we did for the client was to automate their ETL processes using Fabric Data pipelines.

Most of the ETL processes involved integrating with third party REST APIs, getting data from the APIs, performing some transformations and aggregations, and finally storing them in either a Lakehouse or a Warehouse or both.

We utilised Python Notebooks to perform this logic with Pyspark, and utilised Data pipelines to orchestrate and schedule the ETL processes.

Business Insights with PowerBI

For most of the data we imported into Fabric, the end goal was a PowerBI report that business users could use to draw meaningful business insights and make key business decisions from these Insights.

We performed dimensional modelling on the raw data we imported, to convert it into facts and dimensions following the star schema, creating meaningful measures and computed columns in the dimensional model, then used the dimensional model to create powerful PowerBI reports, with detailed and intuitive visuals for ease of interpretation for the end business user.

I ran and managed the delivery of a few of these processes with a small sub-team under me, to see the process from end to end, thus from ETL to insights.

Technology stack used across projects:
- C# .NET
- Azure Cloud: Functions, Logic Apps, Blob Storage
- ARM templates (Azure Cloud IaC)
- Microsoft Fabric Platform
- MSSQL (T-SQL)
- Python (Synapse Notebooks in Fabric)
- Fabric Data pipelines
- Power BI
- SSIS

Duties within the team:
- Designing and implementing data engineering solutions on Microsoft Fabric and on Premises with SSIS
- Data analytics and reporting with PowerBI reports
- Participating in design sessions, breakdown and estimation of tasks.
- SQL Server Stored procedures and Functions
- Developing IaC with ARM templates to easily re-produce our infrastructure
- Azure Functions in C# .NET
- Participating in agile rituals and facilitating scrums, retros, PI planning session.
- Technical presentations to the team on work completed.
- Stakeholder interactions
- Running and managing the delivery of sub-projects with a sub-team.

Betway Africa, Johannesburg

AUG 2025 - PRESENT

Senior AI Engineer

AUG 2025 - PRESENT

In this fast paced role, I hit the ground running, working on several projects in parallel, some of which have an iteration already in production. A highlight of the projects in question are, without exhausting the full list, as below:

AML, Fraud and Responsible Gambling

The biggest project, spanning across three departments in the organization. The task was to implement machine learning based solutions to help with detecting Fraud, potential money laundering activities and patterns, as well as potential gambling abuse behavior. The business had a lot of data but none of it was prepared for machine learning, and there was no already prepared feature store to train algorithms on or ground truth from which training features for supervised learning can be built.

The strategy employed was to utilize unsupervised learning first, utilizing algorithms like K-Means clustering, Isolation Forest and Gaussian Mixture models to categorize gamers into 3 categories of “high risk”, “low risk” and “medium risk” in terms of Fraud, AML and Responsible gambling. We then relied on business users to review, through a web portal we built, the high risk cases and tag them accordingly. This way we would be building our own ground truth for use in supervised learning later.

Through the project, we worked closely with the data engineering team, and business users to craft out new features, and utilize them for model training and prediction.

The technology stack used was:

- Microsoft Fabric Notebooks, Data pipelines, Warehouses and Lakehouses
- Python
- Azure DevOps
- NextJs (ReactJs)
- Azure AppService
- Azure Static Web Apps
- Docker

RAG system

On this project, we built a Retrieval Augmented Generation (RAG) interface that users could utilize to interact with business knowledge via a chatbot like user interface. The business knowledge was integrated from various sources such as SharePoint, OneDrive and Local Disk. For SharePoint and one drive, we implemented a webhook that would listen for file changes and update the vector store accordingly so that the latest, most accurate and most recent information was made available to the chatbot users in near real time.

The technology stack used was:

- Azure Search Service
- Azure App Service
- Microsoft Graph API
- SharePoint and OneDrive
- NextJS (React)
- Python (Fast API, TKinter)
- Docker

OCR powered Expense Tracking System

On this project, the business was struggling with expense capturing, claims processing and record keeping of expenses. They had been considering buying an already existing system. We then utilized our AI product solutioning skills to craft up a solution in-house.

We utilized Azure Document Intelligence studio to train a model that can extract key fields from receipts. We then built a NextJs and Python FastAPI Web solution around this model where users can either open their camera and capture a receipt or upload an already captured receipt from their file storage, which will then be parsed by the model and key information extracted and persisted in a Fabric Lakehouse, along with the URL of the blob which we utilized to display in the application when users would review their expenses. We then built a Power BI dashboard on top of the data for ease of analytics.

Technology stack used:

- Azure Document Intelligence
- NextJS (React)
- Python (Fast API)
- Microsoft Fabric Lakehouse and Warehouse
- PowerBI