

RAJAN KUMAR

Toronto, ON • (437) 986-3952 • rajankumar656@gmail.com • linkedin.com/in/rajankumar95 • github.com/imrajankumar95

SUMMARY

Data and cloud engineering professional with 17 months of enterprise experience at TD Bank, leading the migration of ~3 million customer records across a multi-phase legacy CRM to Salesforce pipeline. Certified Azure Data Engineer (DP-203), DataExpert.io Combined Excellence graduate (Data Engineering + Analytical Engineering), and DataCamp Data Engineer in Python. Currently completing a Cloud Computing diploma at George Brown College, targeting co-op and entry-level Cloud Data Engineering roles across Canada.

TECHNICAL SKILLS

Data & ETL: Azure Data Factory, Azure SQL Database, Salesforce CRM, Python ETL Automation, SQL Pipelines, Data Modeling

Cloud Platforms: AWS (EC2, S3, VPC, IAM, Lambda, RDS, ECS, EKS, ECR, Bedrock, CloudFormation) | Azure (VM, App Service, Blob, VNet, Data Factory, Azure SQL, Azure AD, Azure DevOps)

DevOps & IaC: Terraform, Docker, Docker Compose, Kubernetes (K3s + AWS EKS), GitHub Actions, CI/CD

Monitoring & Observability: Prometheus, Grafana, Node Exporter, AlertManager, Azure Monitor (building)

Programming: Python, SQL, Bash, PowerShell

Tools & OS: Git, GitHub, VS Code, Jira, Agile/Scrum, Draw.io, Linux (Ubuntu, RHEL), Windows Server | IAM, RBAC, VPC Architecture, Security Groups

CERTIFICATIONS

Completed: Azure Data Engineer Associate DP-203 | Microsoft • Data Engineer in Python | DataCamp | Jan 2025 • Agentic AI on AWS | BESA | 2026
Free Data Engineering Bootcamp | DataExpert.io • Combined Excellence — Data Engineering + Analytical Engineering | DataExpert.io

In Progress: Azure Administrator Associate (AZ-104) | Microsoft • AWS Solutions Architect Associate | AWS - both expected 2026

PROFESSIONAL EXPERIENCE

IT Consultant | TD Bank | Toronto, ON

May 2022 – Oct 2023

CRM Data Migration & Data Engineering — Legacy Enterprise CRM to Salesforce on Azure

- Led migration of ~3 million customer records from a legacy enterprise CRM to Salesforce on Azure across 4 business units; designed and executed a multi-phase data pipeline delivering zero-downtime cutover with no data loss.
- Built Python and SQL ETL automation scripts to extract, transform, and load large volumes of customer data — reducing manual processing effort by 40% and eliminating recurring data quality errors.
- Configured Azure Data Factory and Azure SQL Database as the core pipeline infrastructure; enforced RBAC and audit logging to maintain data security and regulatory compliance throughout all project phases.
- Executed UAT data validation test plans, resolved critical data integrity issues prior to go-live; authored migration runbooks and presented weekly pipeline status updates to business and compliance stakeholders.
- Operated within an Agile/Scrum framework with 2-week sprints; tracked data migration milestones and cross-team dependencies in Jira.

PROFESSIONAL DEVELOPMENT

Data Engineering Upskilling | Self-Directed

Oct 2023 – Dec 2025

Completed the DataExpert.io Data Engineering Bootcamp and DataCamp Data Engineer in Python track; earned Microsoft Azure Data Engineer Associate (DP-203) certification while transitioning into cloud and data engineering.

EDUCATION

Ontario College Diploma — Cloud Computing Technology | George Brown College, Toronto, ON

Jan 2026 – Dec 2026

Program includes mandatory industry co-op placement (Fall 2026) | Coursework: Cloud Infrastructure & Virtualization, Linux Administration, Cybersecurity Fundamentals, DevOps, Database Management

ACADEMIC PROJECTS

The Migration Arc | Build & Deploy | George Brown College

Mar 2026 – Present

- Provisioned an Ubuntu VM using Vagrant, installed and configured Nginx, and deployed a static web application locally — establishing a reproducible on-premises baseline for the full migration arc.
- Containerized the application by writing a Dockerfile, building and testing locally, then pushing to Docker Hub and AWS ECR; implemented IAM access policies to govern image pull permissions across environments.
- Provisioned cloud infrastructure using Terraform, automating VM and container cluster setup; deployed the containerized app to AWS ECS — eliminating manual configuration drift across environments.
- Extended to a Kubernetes-native architecture: installed K3s locally, authored Deployment and Service YAML manifests, then deployed identical manifests to AWS EKS — zero changes across environments. Application is end-to-end observed by the monitoring stack below.

Infrastructure Monitoring | Observe | George Brown College

Mar 2026 – Present

- Designed and deployed a full observability stack using Docker Compose — Prometheus for metrics collection, Grafana for visualization, and Node Exporter for host-level metrics (CPU, memory, disk, network) — monitoring the app deployed in the Migration Arc project above.
- Built custom Grafana dashboards and configured AlertManager with Slack and email notification channels for threshold-based alerting, replicating production-grade monitoring practices.
- Deployed the identical stack to an AWS cloud VM, validating full portability across environments and completing an end-to-end build → deploy → observe pipeline.