

Vinoth Venkatesan

(972)-802-2454 | vinothmail20@gmail.com | McKinney, TX 75070 (Open to Relocate)

Experience

Data Engineer | The Home Depot

Mar 2024 – Present

- Architected **Snowflake star-schema** and **event models** powering **AI-driven analytics** across **600+ branches**, enabling smarter operational decisions at scale.
- Built a hierarchical **compliance rule engine (Matillion + Snowflake)** validating **5M+ eCommerce products/day** contributing to **\$1M+ in digital revenue**.
- Delivered a **real-time Route Planner** using **Snowflake delta detection**, syncing only modified orders and streamlining logistics operations.
- Built a **Snowflake-native observability framework (freshness, volume, schema drift)** catching anomalies before nightly **ELT** runs and boosting data trust.
- Launched a **Streamlit + Cortex AI** cost-monitoring app with **2-sigma anomaly detection** and **object-level credit attribution** — driving proactive cost savings.

Data Engineer | NerdWallet

Sep 2023 – Feb 2024

- Cut **Snowflake compute spend** by **\$15K/month** via **warehouse right-sizing, workload isolation, and auto-suspend tuning**; mentored team to scale practices org-wide.
- Re-architected **legacy ETL** with **Airflow**, reducing batch runtime by **80% (8 hrs → 45 min)** for faster business insights.
- Tuned **Airflow** tasks to leverage **Spot Instances**, saving an additional **\$2K/month** in compute.

Data Engineer | Affirm

Jul 2021 – Aug 2023

- Migrated **55+ LTV modeling pipelines** to **Snowflake dbt models** for **Marketing, Growth Analytics & Ops Analytics** enabling trusted **attribution insights** that drive marketing investment decisions.
- Built **DBT models (staging - intermediate - mart)** powering the data science team's **MMM model** decomposing sales into baseline vs. marketing-driven impact to guide spend.
- Automated marketing reporting via **Databricks–Snowflake–Google Sheets** with **scheduled refreshes**, saving business teams **2–4 hours per cycle**.
- Built a dedicated **Snowflake consumption layer** accelerating **Tableau-to-Looker Studio migration** and cutting dashboard development time.

Hadoop Developer | State Farm

Apr 2020 – Jul 2021

- Engineered **Spark (Scala/Python) ETL** jobs parsing multi-format event logs (**JSON/CSV/text**), accelerating downstream analytics.
- Integrated **Kafka** as the **real-time ingestion** layer for high-volume distributed event streams into **Hadoop**, with tuned **topics, partitions, and replication** for **fault tolerance**.
- Detected **25+ fraudulent activities** via **web-log keyword filtering**; delivered warehouse solutions improving transaction system performance by **45–65%**.

Associate Consultant | British Telecom

Mar 2019 – Mar 2020

- Improved **SQL query performance** by **70%** via **strategic indexing** and removal of inefficient **cursors/temp DB** usage.
- Automated **quarterly pricing pipelines (staging → dim → fact)**, eliminating manual effort and ensuring accuracy across all telecom products.

Oracle Developer | Humana

May 2017 – Feb 2019

- Tuned **80+ slow SQL queries** in **SSIS ETL** packages using **explain plans** and **Oracle hints**, cutting runtime by **60%**.
- Identified **\$1M+** in annual over/underpayments in medical claims through **DataMart** analysis support.

Product Specialist | IVY CPG

Oct 2015 – May 2017

- Built custom **SSRS reports** improving department workflow by **45%**; resolved **30+ production tickets** ensuring data accuracy.

Technical Skills

Programming Languages: Scala and Python

Hadoop Components: HDFS, MapReduce, Yarn, Spark, Pyspark and Cloudera Manager

Database: Oracle, MS SQL Server, Hive and Snowflake

ETL: Matillion, DBT, Fivetran and Databricks

Cloud Technologies: AWS – Glue, Redshift, Athena and S3

Workflow Engines: Oozie and Airflow

Visualization: Looker and Streamlit

Messaging & Streaming: Apache Kafka

Accomplishments & Certifications

- Microsoft Certified: Azure Data Fundamental
- DP-700: Microsoft Fabric Data Engineer Associate
- SnowPro Core Certification
- SnowPro Advanced: Architect Certification
- Fivetran Technical Certification
- Databricks Certified Data Engineer Associate
- Oracle Certified: Oracle PL/SQL Developer Certified Associate

Education

Master of Computer Application - MCA

Jul. 2011 – Aug. 2014