# Bharath Surampudi

Data Engineer | Real-Time Payments Data Platforms | Snowflake · dbt · Airflow · Kafka/Flink · AWS
0410 638 861 | bharathsurampudi@gmail.com | linkedin.com/in/bharathsurampudi | github.com/bharathsurampudi | bharathsurampudi.com
Sydney, Australia · Permanent Resident

## Summary

Data Engineer with 5+ years of experience, including 5 years in regulated payments, currently at Australian Payments Plus running and improving the data platform behind NPP, eftpos and ConnectID. Run 200+ production Airflow DAGs on Snowflake and dbt, resolving 5–20 production incidents a month, and build AI-assisted engineering workflows with custom agent skills for spec-driven development. Previously 4+ years at Mastercard on high-availability, PCI-regulated payment data pipelines. AWS Certified Data Engineer and SnowPro Core certified, with a Master's in IT (AI & Database Systems) from UNSW.

## Experience

**Australian Payments Plus (AP+)** Sydney, Australia
*Data Engineer* *Mar 2026 – Present*

- **Owned Snowflake RBAC end-to-end**, from proof of concept to production solution, defining the role hierarchy and rolling out least-privilege access across 5 projects handling sensitive payments data.
- Run BAU for **200+ production Airflow DAGs** across NPP, eftpos and ConnectID on Snowflake and dbt, resolving **5–20 production incidents a month** through triage, root-cause fixes and reruns.
- Reduced BAU toil with pipeline health monitoring for Airflow and GitLab CI, with downstream Snowflake lineage impact tracking, and built a Dynatrace audit-log pipeline (S3 → Snowpipe → Snowflake → dbt).
- Developed **custom agent skills for Snowflake Cortex Code and Kiro** that encode team standards (dbt naming and required tests, Airflow DAG conventions, Snowflake user onboarding and offboarding) into a spec-driven development workflow, and built an LLM-maintained documentation system for DAGs and dbt models.
- Implemented PIM just-in-time access for privileged admin roles, removing standing admin privileges.
- Contributed to a real-time transaction data platform streaming payment events via Confluent Kafka and Apache Flink into Snowflake.

**Mastercard** Sydney, Australia
*Software Development Engineer II (Data & Payment Gateway Systems)* *Nov 2021 – Feb 2026*

- Reduced mean time to resolution (MTTR) for data transfer incidents by ∼40% through improved observability using Splunk dashboards and NiFi provenance log parsing.
- Delivered cross-region data migration pipelines moving multi-terabyte datasets securely across global regions in PCI-regulated environments.
- Architected and maintained high-volume ingestion pipelines routing payment transaction data across multiple regions using Apache NiFi and Java, achieving 99.99% availability under strict SLAs.
- Enforced schema integrity using consumer-driven contract testing (Spring Cloud Contract) across 5+ producer and consumer services, safeguarding financial and regulatory reporting accuracy.
- Partnered with Principal Architects on ETL modernisation, decoupling monolithic pipelines into modular cloud-aligned ELT patterns, improving throughput and maintainability.
- Developed automated data quality and validation workflows integrated into CI/CD pipelines to maintain consistent compliance standards.

**Neau Collective** Sydney, Australia
*Software Engineer* *Mar 2021 – Nov 2021*

- Built Python ingestion pipelines extracting and normalising data from APIs including Shopify and marketing platforms, reducing manual reporting effort by ∼70%.
- Consolidated multi-source business data (Sales, Marketing, Accounting) into unified analytical datasets, enabling near-daily reporting for downstream BI.

## Projects

**Event-Driven Real-Time Fraud Detection Platform** | *AWS Kinesis, Lambda, DynamoDB, Terraform, Athena*

- Serverless streaming pipeline for near real-time fraud scoring, with idempotency (Lambda Powertools) guaranteeing exactly-once processing of duplicate events.
- DynamoDB models for SCD Type 2 location history and velocity rules using atomic counters; Pydantic schema validation and an S3 cold path for audit via Athena.

**Medallion Architecture Data Platform** | *AWS S3, Glue, Redshift, dbt, Airflow, Terraform*

- Bronze-Silver-Gold architecture transforming raw JSON events into dimensional models with dbt, with dbt tests and Great Expectations checks at every layer.
- Airflow-orchestrated ELT deployed via GitHub Actions, with Redshift and S3 infrastructure provisioned in Terraform.

## Technical Skills

**Languages**: Python, SQL, Java
**Cloud & Data Platforms**: Snowflake (Snowpipe, Time Travel, RBAC), AWS (S3, Glue, Lambda, Kinesis, Athena, Redshift, DynamoDB), Apache NiFi
**Streaming & Processing**: Confluent Kafka, Apache Flink, AWS Kinesis, Event-driven architectures, Idempotency patterns
**Orchestration & Transformation**: Airflow, dbt, Pandas, PySpark
**Security & Governance**: RBAC, PIM / just-in-time access, Least privilege, PCI-regulated data, Data contracts
**DevOps & Observability**: Terraform, GitLab CI, GitHub Actions, Docker, Splunk, Dynatrace
**AI-Assisted Engineering**: Agent skill development, Spec-driven development, Snowflake Cortex Code, Kiro, Claude Code
**Data Architecture**: Medallion architecture, Dimensional modelling, Schema evolution, Data lineage

## Certifications

AWS Certified Data Engineer – Associate (DEA-C01) · SnowPro Core Certification (COF-C02) · Databricks Fundamentals Accreditation

## Education

**University of New South Wales** — Sydney, Australia
*Master of Information Technology – Dual Specialisation: AI & Database Systems* — *2019 – 2021*

**Vellore Institute of Technology** — India
*Bachelor of Technology in Computer Science* — *2014 – 2018*