

Tejas Patil

Data/ Analytics Engineering

tejas0412@gmail.com | (437) 422-4636 | [LinkedIn](#) | [GitHub](#)

SUMMARY

Data & Analytics Engineer (AI Certified) with 8+ years of experience, bridging Business Intelligence and Data Engineering by building scalable ETL/ ELT pipelines, dimensional models, and cloud data warehouses (Snowflake, Redshift, Azure). Hands on with SQL, Python, Airflow, PySpark, and AWS/Azure, turning raw data into trusted datasets that power executive dashboards, A/B tests, machine learning models influencing \$100M+ in product & operational decisions.

RELEVANT SKILLS

SQL, Python(NumPy, Pandas, PySpark), R, ETL/ ELT, Data Pipelines, Medallion Architecture, Data Modeling (Star/Snowflake schema), Data Warehousing, dbt, Databricks, Orchestration (Apache Airflow), Data Connectors, AWS (Redshift, S3, Glue, Lambda), Snowflake, Azure (Data Warehouse/Synapse), Google Cloud Platform (GCP), MS SQL Server, PostgreSQL, MySQL, Tableau, Quicksight, Looker (LookML), PowerBI, Preset, A/B Testing, Experimentation, Agile, Data Governance & Data Security, Performance Optimization, CI/ CD, Version Control, Ethics and Compliance (PII data)

EXPERIENCE

TikTok USDS

Data Engineer, User Support Operations (USO)

March 2025 – Sept 2025

Led data engineering and data platform design for User Support Operations, building scalable data infrastructure for analytics, A/B testing & experimentation, enabling data-driven capacity planning and operational cost reduction

- Designed and implemented Snowflake medallion architecture (bronze/silver/gold) using Python & SQL to ingest & transform support ticket and enforcement data with dbt and GitHub Actions CI/CD. Delivered quality checked data marts that informed workforce planning, projecting 180K monthly savings.
- Orchestrated Snowflake dimensional models via Airflow DAGs, building analysis ready datasets for A/B testing agent coaching tools. Enabled ops team to measure impact, correlating with 24% faster handle times and 18-point CSAT gain.
- Built PySpark transformation pipelines with SLA and data quality monitoring to keep key support dashboards within 15 minute freshness, giving leaders proactive visibility and helping cut escalations by 31%.

Amazon

Business Intelligence Engineer II, Fulfillment By Amazon (FBA)

May 2022 – Feb 2025

Spearheaded business intelligence and data engineering for FBA shipment visibility, building AWS based data pipelines, analytics ready datasets, and executive dashboards for seller facing features and internal operations

- Built shipment data pipelines using AWS Glue, S3, and Redshift to power Delivery Windows feature, improving arrival ETA accuracy by 45% and supporting 130M/year in estimated operational savings.
- Developed Python based ETL jobs in AWS Glue to cleanse and enrich large shipment datasets, boosting tracking ID accuracy by 40% and enabling automated seller coaching system.
- Modeled complex shipping pricing logic in Redshift using advanced SQL window functions and Git-based version control, enabling pricing analyses that informed rate changes and lifted seller retention to 90% for non-Amazon carriers.
- Architected reporting datasets on AWS, organizing seller and shipment data in S3 and Redshift with efficient partitioning and Redshift sort/distribution keys, slashing query times for critical seller tools by 30%.
- Utilized SQL, Python, and Looker (LookML) to support A/B tests on seller tools defining metrics and dashboards that helped product teams validate hypotheses and speed adoption of new shipment visibility features.

Amazon

Business Intelligence Engineer II, Abuse/ Fraud Prevention

July 2021 – May 2022

Led analytics and data solutions for competitor abuse prevention against Amazon sellers to reduce overall abuse

- Redesigned seller-on-seller abuse detection criteria using Redshift SQL based rules and reporting, capturing 30% more abusive activity across key marketplaces.
- Engineered feature & evaluation datasets in Redshift for ML driven auto cancel of high risk orders that contributed to >50% reduction in targeted abuse patterns.
- Built risk scoring model in SQL to prioritize top offender investigations. Increased identification of malicious actors 20% by surfacing high risk sellers to enforcement teams.

- Performed deep dive analyses on seller escalations with SQL, Python, and Amazon Quicksight, quantifying risk and helping protect an estimated 1.2B from abusive behavior.
- Collaborated with data science & engineering teams to design analytical & feature-style tables in Redshift data warehouse, persisted to S3 for SageMaker models used for fraud models and automated order cancellation, driving 10% more new abuse case submissions by enforcement teams.

Cadence Design Systems

Data Analyst, People/ HR Analytics

Nov 2019 – July 2021

Led analytics for People/ HR and Finance, building data pipelines, dashboards, forecasts, and reporting workflows that improved visibility into headcount, hiring pipeline, and workforce metrics

- Built an automated Tableau staffing dashboard combining requisition, pipeline and hire data, giving recruiters real time visibility and cutting manual reporting effort by 75%.
- Developed a headcount forecasting model in Python using Holt winters time series methods, helping HR leaders plan hiring and diversity targets across business units and track progress over time.
- Built automated Azure Data Factory pipeline migrating HR and Finance data from MySQL and Oracle into Azure Synapse (Azure data warehouse), cutting manual report prep time by 80%.
- Delivered valuable insights on headcount, attrition, and compensation using SQL, Python, Tableau, and advanced Excel, helping leadership build retention programs targeting key attrition drivers.

VMware (via Decision Minds)

Data Analyst

Oct 2018 – Nov 2019

Led data engineering and dashboard development of multiple tracks including people, bookings, finance, and partners

- Designed and built an FP&A dashboard using SQL and Tableau to combine bookings, revenue, and expense data, automating MBR, and QBR manual Excel reporting and improving finance workflow efficiency by 70%.
- Refactored the Partner Portfolio dashboard by streamlining data loading queries, calculations, filters, and layout, improving dashboard load and interaction performance by 80%.
- Reduced data retrieval time by 50% by optimizing data flow architecture, and implementing ETL pipelines to consolidate data from SAP HANA, PostgreSQL, HAWQ, Oracle, and MySQL into Amazon Redshift via S3 staging.
- Analyzed bookings, finance, HR, and partner data, with SQL across SAP HANA, PostgreSQL, HAWQ, Oracle, and MySQL), generating insights that helped senior leadership identify cost saving and revenue opportunities.

EDUCATION

Santa Clara University

Master of Science Degree in Information Systems

University of Pune

Bachelor's Degree in Computer Engineering

Scrum Alliance

Certification in Certified Scrum Product Owner (CSPO)

Brainstation

Certification in Artificial Intelligence (AIC)