

Phani Gudipally

Manager, Data Engineering

(408) 650-9212

pgudipally@gmail.com

Seattle, WA

[LinkedIn](#)

Summary/Objective

Data Engineering Manager with 12+ years of experience leading data-driven initiatives and over 6 years of managing global engineering teams. Proven expertise in data modeling, cloud infrastructure, modern data stacks, and delivering scalable analytics platforms. Strong technical leadership, cross-functional collaboration, and a passion for building resilient data ecosystems that drive business value.

Skills

- **Leadership & Strategy:** Team building, mentorship, and roadmap ownership. Strategic alignment with business goals. Manage remote and global teams
- **Data Modeling & Warehousing:** Scalable data architecture, dimensional modeling. Hands-on with Snowflake, Redshift, BigQuery
- **ETL/ELT & Orchestration:** Pipeline development using Airflow and DBT
- **AI & LLM Tools** OpenAI, Perplexity APIs, LangChain, MCP servers, automated SQL generation, AI agent development
- **Programming & Scripting:** Advanced Python and SQL for data processing, automation, and analytics.
- **Analytics Engineering:** Metric standardization, Reusable data marts, and enablement of self-serve analytics across domains like sales, product, and operations.
- **Data Governance:** Data quality, metadata management, lineage, and compliance frameworks.
- **Visualization & BI:** Self-serve analytics platforms. Skilled in Looker, Tableau, Superset.

Work Experience

MANAGER, DATA ENGINEERING, SHOPIFY

Revenue Data Engineering | Seattle | Jan 2024 – Current

- Built Shopify's first Revenue Data Engineering team from the ground up, supporting a \$5+B revenue pipeline
- Led technical architecture and cross-functional design for Shopify's Data Platform, including evaluating and migrating analytics platforms (BigQuery, Snowflake, Databricks).
- Established the technical vision for the Revenue Data Engineering organization; architected and built the revenue data platform from scratch using DBT, Airflow, and the full suite of cloud data services.
- Collaborated with Sales, Growth, Product, Marketing, MarTech and Data Science to design and deliver reliable data products and reporting solutions.
- Developed the technical roadmap for the Revenue Data Engineering team, setting a clear vision for platform growth and innovation. Established Org strategy to ensure all initiatives were tightly aligned with revenue operations and business objectives
- Spearheaded the design of AI-native data models and intelligent MCP servers, optimizing schema architecture and context management for LLM performance. Streamlined self-serve analytics and automated data enrichment, resulting in a 40% increase in adoption among non-technical users and a 30% reduction in manual reporting requests.
- Developed AI agents using OpenAI and Perplexity APIs to enrich internal sales lead data by extracting contact and contextual information from lead sources, enhancing lead qualification accuracy and significantly boosting rep conversion rates.
- Mentored engineers, drove best practices in testing, modeling, and automation, and fostered a culture of accountability and excellence.
- Introduced agile practices and efficient prioritization processes to deliver critical initiatives on time and within scope. Drove operational efficiency in sales and onboarding workflows by delivering trusted, automated data pipelines and reporting layers.
- Engineered AI-ready testing frameworks in dbt for automated model validation and documentation. Integrated with CI/CD to ensure data quality and trust in LLM-powered analytics workflows.

MANAGER, DATA ENGINEERING, AMAZON

Product Compliance, Compensation | Seattle | Mar 2020 – Sept 2023

- Led the architecture and management of Amazon's large-scale, cloud-native data infrastructure supporting 2,000+ active users, leveraging AWS Data Lake, Redshift, EMR (Spark, Hadoop), and S3 for high-throughput, resilient analytics and storage.
- Built and optimized data platforms to process petabyte-scale data with Redshift, Kinesis, Glue, Lambda, DynamoDB, and Athena, ensuring high performance, low latency, and robust security
- Designed and implemented data products and solutions that enabled rapid product compliance investigations, reducing investigation cycle times by 80% through automated pipelines and self-service analytics.
- Achieved major infrastructure cost reduction goals by optimizing storage and compute (resource tagging, tiered storage, auto-scaling, implementing proactive observability, analyzing usage trends, and building systems to keep storage and compute costs in check, resulting in significant operational savings.
- Set and communicated the long-term strategic vision for data infrastructure, working closely with senior leaders in Product, Engineering, and Business Operations to align on priorities and anticipate future needs. Ensured the data platform could scale and adapt to evolving business goals. Regularly led roadmap sessions and strategy reviews to drive cross-functional alignment and shared accountability.
- Mentored and coached a high-performing team of 6 engineers, driving a collaborative, high-ownership culture, 95% retention, and a 20% boost in team productivity through targeted mentorship and skill development programs.
- Managed engineering teams through solution design, development, and deployment of data lakes and real-time pipelines.
- Set long-term strategic vision and collaborated with senior stakeholders across product and engineering to align goals.

SENIOR DATA ENGINEER, AMAZON

Alexa AI, Amazon Devices, Product Compliance | Seattle | 2015 – 2020

- Led the design, implementation, and successful delivery of large-scale, critical data solutions involving AWS services including Redshift, Kinesis, S3, EC2, RDS, Glue (Spark), EMR, Lambda etc.
- Build cross-functional relationships with Data Scientists, Product Managers and Software Engineers to understand data needs and help them adopt best practices in reporting and analysis
- Build reliable ETLs/ELTs to take data from various sources, data lakes and supporting data sources and develop optimized data model for analytics and reporting.
- Consistently improved reporting and analysis processes, pushed towards automating or simplifying self-service support for internal stakeholders

SENIOR DATA ANALYST, PAYPAL -LATENTVIEW ANALYTICS.

Digital Marketing Analytics | San Jose | 2010 – 2014

- Gather business requirements, translate business needs into detailed analytical requirements, and ensure delivery timelines are met and that the deliverables meet the slated business objectives
- Design and analyze the effectiveness of campaigns to help the client increase their market share and revenue
- Prepare high level dashboards for the marketing heads using Visualization tools like Tableau

Education

BITS - Pilani

B.E. (Hons.) Mechanical Engineering, 2006-10

Rajasthan, India