

Gokul Nandakumar

goku.careers@gmail.com | Portfolio | github.com/Goku007007 | LinkedIn | +1 (773) 930 2964

Skills

Languages: Python, JavaScript, SQL, Shell Scripting, Java **Frameworks & Tools:** PHP(Laravel), REACT, TypeScript, NodeJS, FastAPI, Flask, Docker, Kubernetes, Terraform, AWS, Azure (ADLS, Synapse), Git, GitHub Actions, Airflow, Spark **Databases:** PostgreSQL, MySQL, DynamoDB, Redis, Cassandra **Concepts:** REST API Design, Microservices Architecture, CI/CD, System Design, Distributed Systems, Agile/Scrum **Certifications:** AWS Certified Solutions Architect - Associate **AI - Assisted Dev & API:** Cursor, Claude Code, ChatGPT-4o, Cursor, Antigravity, LLM and Vid Gen APIs

Software Engineer, ChiEAC - Chicago, IL (Remote) Jun 2025 - Present

- Designed and implemented scalable backend services using Python and Airflow, reducing data processing time by 66% (6h → 2h) through concurrent task execution and optimized resource allocation
- Architected distributed data systems on Azure, unifying 5TB+ of disparate data into a centralized lakehouse to enable real-time analytics for city wide transit systems.
- Built automated reporting microservices using Python and SQL, eliminating manual report generation and standardizing 12 core KPIs for executive leadership.
- Improved system reliability and auditability by designing reproducible pipelines using Delta Lake and Medallion Architecture, ensuring data consistency across civic and transit analytics workflows.

Software Engineer, ISQL Global - Remote Oct 2022 - Jul 2023

- Re-architected legacy REST APIs using Python (FastAPI) to optimize data retrieval, cutting redundant backend calls by 30% and reducing payload size, resulting in a 30% faster load time for the reporting feature.
- Implemented automated testing and CI/CD pipelines using GitHub Actions, ensuring 99.9% data quality reliability and enabling rapid deployment of schema changes.
- Optimized backend data processing pipelines using PySpark and Airflow, reducing data latency by 87% (2h → 15m)
- Enhanced database query performance by 35% and reduced compute costs by 15% on Redshift/Databricks by optimizing SQL join logic and implementing effective caching strategies.

Software Engineer, Out In Dreams - Remote Sep 2021 - Oct 2022

- Developed an automated payment processing system using Python, reducing manual finance operations by 20+ hours/week and eliminating inconsistent payment errors for the referral-based program.
- Integrated third party APIs (Google Ads) into internal data warehousing solutions, creating a unified view of ad spend and ROI for cross-functional teams.
- Engineered data models in Snowflake to support campaign performance tracking, reducing reporting turnaround time by 40% (5 days → 3 days).
- Analyzed \$20k+/month in ad spend data to identify optimization opportunities, resulting in a 12% improvement in campaign ROI through algorithmic budget reallocation.

Projects

National EV Charging Infrastructure Analytics [Github](#)

- Developed a data ingestion and analytics platform for 83k+ EV charging station records, utilizing (GCS) and BigQuery
- Deployed an interactive dashboard to visualize network coverage gaps, processing large scale geospatial data to identify underserved regions.

GPU Cluster Telemetry & Observability System [Github](#)

- Built a scalable observability platform using Python and DuckDB to monitor GPU cluster performance, ingesting 10M+ trace rows for granular utilization tracking.
- Implemented an IsolationForest machine learning model to detect abnormal cluster behavior, reducing false alerts

Education

Illinois Institute of Technology - Master of Information Technology

SREC - Bachelor of Computer Science and Engineering