

JUSTIN JENISH

Anaheim, CA | (213) 284-0593 | justin4jenish@gmail.com | [linkedin.com/in/justin-jenish](https://www.linkedin.com/in/justin-jenish) | github.com/pip-install-justin

Solutions Architect for AI & data platforms with hands-on engineering roots — MS in Computer Science (USC) and production experience since 2020 spanning streaming systems (Kafka, Spark), lakehouse architecture (Databricks, Snowflake, Delta Lake), and applied GenAI/RAG across Azure and AWS. Progressed from building and operating pipelines to designing end-to-end architectures and defending them in front of stakeholders — from city procurement teams to sales leadership. Snowflake SnowPro, Microsoft Azure Data Engineer, and Databricks certified.

EXPERIENCE

TSL — *GSA-cleared government services & electronics contractor*

Los Angeles, CA

Lead Full-Stack Engineer & Solutions Architect

Oct 2025 – Present

- Lead software engineering for government RFP responses; delivered live product demos to the City of Los Angeles, translating technical capability into stakeholder-facing value and advancing procurement opportunities.
- Founding engineer of an internal social-platform venture, owning architecture and delivery across frontend (Flutter/Dart), backend, and AI; designed the advertising and revenue-share engine, using blockchain-based settlement records for advertiser-auditable spend.
- Sole software engineer on an autonomous mobile battery-unit prototype (next-generation portable power station), building control software and systems integration for field-deployable energy delivery.

Aipower Inc. — *AI products startup*

Los Angeles, CA

AI Engineer

Jun 2025 – Oct 2025

- Improved chatbot answer precision on exact-term queries (SKUs, model numbers) and cut response latency versus a dense-only baseline by architecting hybrid retrieval (dense semantic + sparse keyword search) for AI sales chatbots deployed in e-commerce guided selling and troubleshooting.
- Led technical discovery with marketing and sales teams, translating customer requirements into AI solution designs and architecture decisions; delivered the contracted product ahead of the department wind-down.
- Built a full-stack, event-driven workflow-management tool coordinating customer service, sales, and inventory across departments.

USC — Viterbi School of Engineering

Los Angeles, CA

Data Engineer (Research)

Sep 2024 – Jun 2025

- Built AMAR, a multi-agent RAG pipeline that turns plain-language questions from non-technical stakeholders into custom analytical reports, with automated retraining and parameter tuning through the Databricks MLflow model lifecycle.
- Cut report turnaround from ~1 hour (manual analyst workflow) to 5 minutes across financial-reporting, healthcare-diagnostics, and customer-support scenarios by orchestrating the pipeline on Databricks and Airflow DAGs, stress-tested at 50M events/day with replayed public clickstream data.

USC — Annenberg School of Journalism & Communication

Los Angeles, CA

Business Analyst / Tech Consultant

May 2023 – May 2024

- Reduced equipment-overdue incidents and A/V and network downtime by building Snowflake pipelines, a data-vault schema, and automated Power BI/Streamlit dashboards; cut response times by streamlining staffing and CRM workflows.
- Developed prescriptive ML models for staffing optimization; led weekly stakeholder reviews and managed server administration with a focus on security.

EARLY CAREER — INDIA (PART-TIME & CONTRACT ROLES DURING UNDERGRADUATE STUDY)

GMG Associates — *Freelance Data Analyst (Finance)*

Bangalore, IN | Jul 2021 – Jul 2022

- Eliminated 10+ hours/day of manual reporting across the team by automating ERP, CRM, and tax-data migration to Azure with Airflow, validation scripts, and unit tests; implemented role-based access for compliance and audit readiness; built Power BI reports and predictive models on Databricks (MLlib, MLflow).

SRM Institute of Science & Technology — *Deep Learning Research Assistant*

Chennai, IN | Jan 2022 – May 2022

- Improved PM2.5 forecasting accuracy — 7.2% lower RMSE, 12.4% higher CSI over baseline models, with 4-day prediction horizons — by engineering a spatiotemporal GNN architecture (GINs + GRUs, PyTorch Geometric) across 184 cities.

Lemonpeak — *staffing firm; client: a major US satellite/streaming TV provider*

Chennai, IN | Dec 2020 – Jun 2021

Data Engineer (Streaming Data)

- Built Lambda-architecture pipelines (Kafka, Spark Streaming, AWS EC2/S3) for subscriber analytics on the client's 5M-daily-user platform; supported ad-ROI tracking, A/B testing, and executive dashboards; mentored 2 junior engineers on testing and warehousing practices (pytest, Great Expectations).

Yatnam Technologies — *Software Development Intern*

Kochi, IN | Feb 2021 – Jun 2021

- Built a UK e-commerce platform (warehouse management + ERP) on REST APIs and ETL pipelines; deployed with Terraform and CI/CD on Azure.

Arcadia Solutions Inc. — Data EngineerRemote | Jan 2020 – Aug 2020

- Cut storage and compute costs ~20% on a Medallion clickstream platform by tuning partitioning, file compaction, and indexing across the Delta Lake storage layer, while supporting predictive sales models and user-behavior analytics.

EDUCATION

USC, Viterbi School of Engineering — MS, Computer ScienceLos Angeles, CA | May 2024

SRM Institute of Science & Technology — BS, Computer Science (Big Data Analytics)Chennai, IN | May 2022

CERTIFICATIONS

[Snowflake SnowPro Core Certification](#) — SnowflakeOct 2024

[Azure Data Engineer Associate \(DP-203\)](#) — MicrosoftDec 2024

[Databricks Certified Data Engineer](#) — DatabricksDec 2024

TECHNICAL SKILLS

- **Languages & Scripting:** Python, SQL, Scala, PySpark, JavaScript, TypeScript, Dart, Bash, HTML/CSS
- **Data Engineering:** Spark, Kafka, Airflow, Flink, dbt, AWS Lambda, Glue, REST APIs, ETL/ELT, Streaming
- **AI & Machine Learning:** Generative AI, LLMs, RAG, Multi-Agent Systems, PyTorch, MLflow, MLLib, Predictive & Prescriptive Modeling, Python data stack (Pandas, NumPy)
- **Warehousing & Databases:** Snowflake, Databricks, Delta Lake, Redshift, PostgreSQL, Cassandra, Hive, S3
- **Cloud & Architecture:** Azure, AWS, GCP, Lakehouse & Medallion Architecture, Data Mesh, Data Modeling, Distributed Systems, System Design
- **Visualization & BI:** Power BI, Tableau, Microsoft Fabric, Streamlit
- **DevOps & Governance:** Linux, Git, Docker, Terraform, Kubernetes (EKS), Helm, Prometheus, CI/CD, Unity Catalog, Collibra, Great Expectations, Data Governance & Quality
- **Collaboration:** Requirements Gathering, Stakeholder Management, Technical Leadership, Agile/JIRA