Jatin Bishla

+91-9996244114 | bishlajatin@gmail.com | linkedin.com/in/jatinbishla | Gurugram, India | *Open to relocation: London · Singapore · EMEA*

Summary

Software engineer with ∼2 years building low-latency, event-driven backends where correctness and cost cannot be traded away. Owns core components of a billing & reconciliation engine processing 1M+ transactions/day – idempotent, exactly-once accounting that stays correct under partial failures, retries, and out-of-order events – fed by Kafka streaming at 10K events/sec. Strong production Python (data tooling, automation, NumPy/Pandas) with a solid data-structures, algorithms, and OOP foundation; actively building C++ systems depth via a limit-order-book matching engine. Drawn to the exact problem class trading infrastructure runs on: ordering guarantees, backpressure, deterministic replay, and p99 latency.

Technical Skills

Languages: C++ (C++17/20), Python, SQL, Bash
Core CS: Data Structures & Algorithms, Object-Oriented Programming, Concurrency, Linux
Systems & Streaming: Low-Latency Event-Driven Architecture, Apache Kafka, Deterministic Replay, Backpressure, p99 Latency
Python & Data: NumPy, Pandas, Matplotlib, Airflow (workflow scheduling)
Infrastructure: AWS (Lambda, SNS, SQS, CloudWatch), Docker, Git, CI/CD

Experience

Acefone — Gurugram, India
Software Engineer – Distributed Systems & Low-Latency Backends — *Sep 2024 – Present*

- Own core components of the billing & reconciliation engine processing 1M+ transactions/day with idempotent, exactly-once accounting – maintaining correct cost attribution under partial failures, retries, and out-of-order events.
- Built event-driven ingestion on Apache Kafka handling 10K events/sec across 10 partitions, with ordering guarantees and backpressure handling to keep the pipeline stable under load.
- Designed fault-tolerant orchestration with deterministic replay, eliminating reconciliation drift across services and making failure recovery reproducible.
- Cut memory footprint of high-volume data export via streaming pagination and columnar (Parquet) storage, improving throughput and cost efficiency on large result sets.
- Built Python tooling for analytics, automated reporting, and post-call insight extraction (speech-to-text + LLM) across 1M+ daily events, replacing manual review and giving stakeholders structured data to act on.

The DataFlow Group — Noida, India
Software Engineering Intern — *Feb 2024 – Jul 2024*

- Refactored serverless workflows on AWS Lambda + Python and built SNS-based real-time alerting, automating monitoring and reducing manual intervention by 40% across production systems.
- Optimized ETL with Apache Airflow to process 100K+ records/day and tuned Redshift warehousing queries, cutting query times by ∼15%; added CloudWatch dashboards/alarms reducing downtime.

Projects

Limit Order Book Matching Engine | *C++17/20* — 2026 (in progress)

- Building a single-threaded, price-time-priority matching engine in modern C++ – integer-tick price representation (no floating point), sequence-numbered events over wall-clock time for deterministic replay, and cache-conscious data structures for fast order insert/cancel on the hot path.

Customer Analytics & Insights Engine | *Python, SQL, Pandas, NumPy, Matplotlib* — Jul 2025

- Built an end-to-end Python analytics platform ingesting 5+ years of customer-level financial data, computing health/risk metrics and generating automated benchmark reports – with structured metric definitions, data-quality checks, and documented methodology.

Education & Achievements

Maharaja Agrasen Institute of Technology (GGSIPU) — New Delhi, India
B.Tech – Electronics & Communication Engineering | *CGPA: 8.49/10.0* — *2020 – 2024*

AI/ML Head – Aerostars Technical Drone Society, MAIT
Led a 25+ member team on AI/ML projects; owned technical direction and communicated complex findings to cross-functional audiences.

Certifications & Contests: AWS Certified Cloud Practitioner · Machine Learning with Python · Google Hash Code 2021 (algorithmic optimization contest)