

Ashish Kadam

+91 8850108480 | ashishkadam2026@gmail.com | linkedin.com/in/ashishkadam89 | github.com/zegulas

SUMMARY

Experienced software engineer with over 10+ years of expertise in building scalable systems, optimizing complex SQL/BigQuery queries, and leading high-performing teams. Passionate about AI integration, chatbot development, and query performance tuning. Adept at identifying technical bottlenecks and delivering innovative, cost-efficient solutions.

EXPERIENCE

Staff SDE

Apr, 2023 – Mar, 2025

GoodGlammGroup

Mumbai

- **Evaluated AWS DocumentDB vs MongoDB Atlas**, benchmarking real-world workloads and uncovering high I/O costs under Standard storage—advised postponing migration until **I/O-Optimized** mode matured
- **Designed advanced rate-limiting for critical APIs** using **Redis** and context-aware identifiers (beyond IP) to prevent abuse without exposing internal logic
- **Built courier performance reports in Data Studio** using **BigQuery** to calculate delays across delivery stages—helping the supply chain team **optimize** courier selection
- **Automated discount abuse detection** via a cron-triggered **Python** script **analyzing** orders-per-minute; sent real-time alerts using **Deadmansnitch** to prevent losses
- **Refactored reporting to an event-driven model**, incrementally aggregating order metrics in real-time and reducing DB load; evaluated **PostgreSQL materialized views** as an alternative
- **Improved MongoDB performance** by **analyzing** slow query logs (**Hatchet**, **MongoAnalyzer**), refining indexes based on ESR rules and validating improvements via Atlas metrics
- **Developed an automated data sync script** to migrate records from prod to pre-prod during off-peak hours with PII redaction for compliance
- **Led BQML ARIMA_PLUS POC to forecast warehouse stock**, **analyzing** seasonality and stockout risks; after cost and accuracy review with Google, advised against production use
- **Built data ingestion pipelines using Cloud Functions + Scheduler**, fetching and transforming third-party API data (courier, warehouse) into **BigQuery**
- **Developed an LLM assistant** that converted natural language to **BigQuery SQL** using prompt engineering and schema context, enabling near-instant insights for leadership
- **Built a GPT-4o-mini chatbot** mapped to 20+ support use cases, integrated with backend APIs for real-time user/order/refund data—automated **~70%** queries, improving CSAT
- **Reduced OpenAI token usage** by building a use-case classifier that narrowed context passed to the LLM, cutting input tokens by **~10×** and reducing API costs
- **Resolved a critical disk alert loop** by **analyzing** log patterns and implementing **logrotate** on the data engineering server—eliminating recurring 100% disk space usage issues
- **Created materialized views over 2TB+ clickstream data in BigQuery** using partitioning and clustering to accelerate queries for data science and analytics teams
- **Optimized ClickHouse performance** by introducing intermediate **materialized** tables and improving aggregation logic—cut compute costs and reduced redundant work
- **Built a GCS → BigQuery ingestion pipeline**, transforming vendor data from external tables and aggregating results using **materialized** views for downstream access
- **Developed real-time anomaly alerts** for order volume and regional **behavior** using **BigQuery** thresholds—flagged potential abuse and operational issues early
- **Identified traffic flagged as attacks as user behaviour misuse**, not a true threat; implemented rate-limiting to reduce abuse while preserving UX for valid users
- **Automated vendor invoice parsing** with a **Python serverless** app that extracted PDFs from emails and converted them into structured Sheets using **Gemini APIs**

Software Engineer

Jun, 2017 – Mar, 2023

MyGlamm

Mumbai

- **Built the first Payment Microservice** alongside a legacy monolith to support full payment and refund flows with essential transaction logging

- **Automated daily order and revenue reports** for each franchise using **Core PHP** and **Mailchimp**, improving reporting accuracy and reducing manual effort
- **Secured social login integration** in an **Express.js** API Gateway by verifying Google/Facebook JWT tokens at the backend, ensuring safe user onboarding
- **Replaced MongoDB reporting with asynchronous BigQuery jobs**, reducing production DB load and accelerating large-scale report generation
- **Cut BigQuery costs by 4×** via partitioning, clustering, and **authorized views**; later switched to flat-rate pricing for predictable cost control
- **Designed real-time config propagation** using **RabbitMQ** fan-out and **Redis** to reflect config updates instantly across all services without redeployments
- **Led a POC for OpenID Connect-based Auth microservice**, evaluated impact on UX/conversion, and advised against deployment due to business concerns
- **Built a Flask service to analyze order patterns and delivery outcomes**, generating risk scores to improve fraud detection and operational decisions
- **Co-designed a performant URL shortener**, **optimizing** key length and evaluated **Bloom filters** as an addon for fast, collision-resistant URL generation
- **Conducted a MongoDB sharding POC**, exploring shard key strategies and **re-sharding** mechanisms to scale high-volume data workloads efficiently

Software Engineer

Mar, 2013 – Jul, 2015

B2 Information Systems Pvt. Ltd.

Mumbai

- **Developed core features** for client-facing projects including Ottawa Dental Society, MProfit, and the BiteBank Admin Panel, contributing to robust and maintainable codebases
- **Progressed from core developer to project lead**, consistently delivering high-quality releases on schedule and mentoring junior team members
- **Led the migration of legacy BiteBank systems** to a modern web framework, streamlining workflows and improving system reliability and maintainability

OPEN SOURCE CONTRIBUTIONS

node-time-sleep

- Python-like `time.sleep(ms)` for Node.js with non-blocking `async/await` support.

json-cache-key-generator

- Deterministically generates a cache key from deeply nested JSON — useful for caching and memoization.

remove-invisible-characters

- Removes ASCII extended characters that are invisible but can break Node.js code — ideal for input sanitization.

SIDE PROJECTS

Global Mute for Google Meet

- Chrome Extension to mute/unmute Google Meet globally using hotkeys — works even outside the browser.
- Built with Manifest V3; live on Chrome Web Store.

Skip Silence

(In Progress)

- Chrome Extension that detects near-silent moments in YouTube videos and speeds up playback to 2× during silence.
- Returns to normal speed once speech resumes — ideal for long lectures.
- Supports real-time audio analysis and customizable fast-forward settings.

Language Learner (German)

- Helps users learn German by showing translated sentences from English, Hindi, or Marathi and assessing pronunciation accuracy.
- Uses 4 transcription methods: SpeechRecognition interface of the Google Web Speech API, OpenAI Whisper (API & FOSS), Azure Speech-to-Text.
- Provides feedback on pronunciation with visual improvement graphs over multiple attempts.

CERTIFICATIONS & COURSES

| | |
|-------------------------------------------------------------------------|----------|
| Building Your Own Database Agent — DeepLearning.AI | Jan 2025 |
| Function-calling and Data Extraction with LLMs — DeepLearning.AI | Jan 2025 |
| Functions, Tools and Agents with LangChain — DeepLearning.AI | Jan 2025 |
| LangChain: Chat with Your Data — DeepLearning.AI | Jan 2025 |
| LangChain for LLM Application Development — DeepLearning.AI | Jan 2025 |
| Finetuning Large Language Models — DeepLearning.AI | Jun 2024 |
| Getting Started with Mistral — DeepLearning.AI | May 2024 |
| Prompt Engineering with Llama 2 & 3 — DeepLearning.AI | Apr 2024 |
| Ultimate AWS Certified Cloud Practitioner CLF-C02 (2025) — Udemy | Mar 2025 |
| JavaScript Algorithms and Data Structures Masterclass — Udemy | Oct 2023 |
| Mastering OpenAI Python APIs: Unleash ChatGPT and GPT4 — Udemy | Aug 2023 |
| Algorithmic Trading using Angel One's SmartAPI — Udemy | Jan 2023 |
| Go: The Complete Developer's Guide (Golang) — Udemy | Nov 2022 |
| Python (Basic) — HackerRank | Sep 2020 |

EDUCATION

| | |
|------------------------------------------------------------------|---------------|
| MMS Operations | 2015 – 2017 |
| <i>Welinkar Institute of Management Development and Research</i> | <i>Mumbai</i> |
| MSc Information Technology | 2010 – 2012 |
| <i>University of Mumbai</i> | <i>Mumbai</i> |
| BSc Computer Science | 2007 – 2010 |
| <i>University of Mumbai</i> | <i>Mumbai</i> |

TECHNICAL SKILLS

Languages & Frameworks: JavaScript, Node.js, Express.js, Flask, Python, Core PHP, CodeIgniter, Laravel
Databases & Storage: MongoDB, PostgreSQL, BigQuery, ClickHouse, AWS DocumentDB, Materialized Views
Cloud & Infrastructure: Google Cloud Functions, Google Cloud Scheduler, AWS Lambda, RabbitMQ, Redis
Data & Analytics: Google Data Studio, Query Performance Tuning, Data Pipelines, Real-Time Data Processing, Data Engineering, Data Visualization, Business Intelligence
AI & Automation: OpenAI GPT-4o-mini, LLM Integration, Prompt Engineering, Customer Support AI Chatbot, Automation
Security & Identity: OpenID Connect, JWT, OAuth 2.0, PII Redaction, RBAC
Architecture & Patterns: Microservices, API Gateway, Event-Driven Architecture, Distributed Systems, Rate Limiting
DevOps & Monitoring: System Monitoring, Logrotate, Hatchet, MongoAnalyzer, Deadmansnitch, Cost Optimization, Anomaly Detection
Other Tools & Concepts: Mailchimp, Postman