

Kumar Shivendu

Search and Distributed (Database) Systems

[✉ mail@kshivendu.dev](mailto:mail@kshivendu.dev) | [🌐 kshivendu.dev](https://kshivendu.dev) | [🔗 KShivendu](https://KShivendu) | [LinkedIn](https://kshivendu) | [Twitter](https://KShivendu_)

Skillset

Languages	Rust	·	Python	·	JavaScript	·	Java	·	Bash	·	C++								
Frameworks	Actix	·	Tokio	·	FastAPI	·	Express.js	·	Django	·	React	·	Material UI	·	Bootstrap	·	Cypress	·	PyTest
Infra	Kubernetes	·	Grafana	·	Loki	·	PostgreSQL	·	Redis	·	Elasticsearch	·	Qdrant (Core engineer)						

Experience

Qdrant: Software Engineer - Core Team

Oct. 2023 - Present

- Qdrant is the **fastest growing vector database with 28k+ stars, 40k+ docker pulls/day, and tens of thousands of production deployments.**
- Conducted performance experiments to investigate our bottlenecks and did comparative benchmarks of Qdrant vs other db/search engines
- **Built the infrastructure for Chaos testing and Continuous benchmarking** which caught 50+ critical bugs and perf degradations early
- **Creator and owner of our cluster manager service** that automatically scales and manages nodes/shards of Qdrant clusters inside k8s.
- Cluster manager **powers features like vertical/horizontal scaling, replication, rebalancing, resharding, multi-zone/region cluster sync**
- Represented Qdrant at Meta Llama Hackathon, managed Qdrant Summer of Code, and delivered 5 talks about vector search and Qdrant
- Dealing with storage and distributed system features & bugs in the qdrant core — consensus, shards, wal, snapshots, shard transfers, etc

SuperTokens (YC 20): Software Engineer

Mar. 2022 - Sep. 2023

- SuperTokens is an authentication devtool with 14k+ Github stars, 2.5M docker pulls, and multiple thousand B2B and B2C SaaS users.
- Architected our most requested **enterprise features - multi-factor authentication (MFA) and Time-based OTP (TOTP)**
- Implemented **authorization (RBAC), multitenancy, MFA, and TOTP features** in our Python and Node.js SDKs

FamPay (YC 19): Software Engineer Intern

Sep. 2021 - Mar. 2022

- FamPay is India's first numberless card and payments app for teenagers with 10M+ app downloads and millions of MAUs.
- Worked on multiple features like UPI payments, KYCs, asynchronous celery batch jobs, and anomaly detection
- Improved Infra and CI by introducing traefik, oauth2-proxy, mypy, parameterized, diff coverage, and developed a tool to debug Redis cache

Software Heritage (Inria Paris): Google Summer of Code Student

May. 2021 - Aug. 2021

- Software Heritage (SWH) is an organization under Inria (Paris) to collect and preserve all of the open-source software that has ever existed.
- Developed **advanced metadata search features on top of 420M+ public software repositories containing 27 Billion+ source code files**
- **Designed a custom search query language** and wrote a **transpiler that converts the Treesitter based AST into Elasticsearch queries**
- Implemented autocompletion for the custom query language in the web interface using WebAssembly, Javascript, CSS, and Django.

Prev. Internships

Feb. 2020 - Jul. 2020

- SDE Intern @ Logy.AI: Developed the appointment booking portal for Colgate MyToothDoc platform using Django (100k+ users)
- ML Intern @ Logy.AI: Improved internal image annotation platform. Implemented Fourier transform based blurry microscopic image detector
- ML Intern @ IIT Mandi: Trained ARIMA and LSTM models to forecast network usage time series data. Our paper was accepted at IEEE

Side projects

Tech talks

- **Not All Tenants Are Equal: Tiered Multitenancy in Qdrant** @ Qdrant Discord Office Hours [Link](#)
- **Qdrant: Internals of a Vector DB** @ Guest Lecture - DBMS Course IIT Bhilai [Link](#)
- **Beyond Filters: Modern Search with Vectors in Django** @ DjangoConUS, Chicago [Link](#)
- **Chaos testing: Breaking a database on purpose** @ PlatformCon (Virtual) [Link](#)
- Improving search relevance using reranking and fusion @ Microsoft Bangalore [Link](#)
- Qdrant: Future of search and beyond @ FOSS United Bangalore [Link](#)

Technical blogs

- I Reverse-Engineered Exa.ai Infrastructure Cost with Napkin Math [Link](#)
- Monitoring my life with Grafana, Prometheus, and InfluxDB [Link](#)
- In-depth look at ACID properties of Postgres [Link](#)

Open Source projects and contributions

- smoldb: **Implemented a distributed database from scratch** with persistence, consensus, and shards etc. [Link](#)
- software heritage: **Received the Google Open Source Peer Bonus** for my continued contributions and **mentored a student in GSoC**
- vscode-advanced-search: VSCode Extension for bulk search/refactor that's more powerful than regex. **Winner in FOSS United Hackathon** [Link](#)
- suggest-ai: Recommendations for everything powered by GPT3 and Google Image Search API [Link](#)
- search-copilot: Chrome extension that extends search on existing websites using LLMs + scraping to give you vector search + filters [Link](#)
- ann-benchmarks: **Did the first ever benchmark of OpenAI embeddings at 1M scale using databases (Qdrant vs Pgvector)**. Creators of Github Copilot, Yahoo Search, Spotify recsys, and Supabase responded to my work on Twitter. This project **accelerated/led implementation of HNSW in Postgres (pgvector)**. The dataset I created has been referred by Cloudflare, Redis, Clickhouse, etc and gets 1K/m downloads [Link](#)

Education

Indian Institute Of Technology (IIT) Bhilai

Bhilai, India

BACHELOR OF TECHNOLOGY IN COMPUTER SCIENCE AND ENGINEERING

Jul. 2018 - May. 2022

CGPA : 8.2/10.0