

# Wahib Sabir Kapdi

Bryan, TX

+1 214 899 2620 | [wahibkapdi2212@gmail.com](mailto:wahibkapdi2212@gmail.com) | [linkedin/wahib-kapdi](https://linkedin.com/in/wahib-kapdi)  
[github/Kanchuking1](https://github.com/Kanchuking1) | [kanchuking1.github.io/myPortfolio](https://kanchuking1.github.io/myPortfolio)

## Education

**Texas A&M University, College Station, TX**

*Master of Computer Science*

- **Coursework:** Algorithms for Big Data, Distributed Systems, Cloud Computing, Parallel Computing

**Aug 2024 - May 2026**

*College Station, TX*

**Birla Institute of Technology and Science, Pilani (BITS Pilani)**

*Bachelors of Engineering, Electrical and Electronics Engineering*

**Jul 2018 - Jul 2022**

*Pilani*

## Professional Experience

**Infineon Technologies**

*Software and Automation Intern*

- Engineered a production-grade C++ DLL exposing stable APIs for hardware communication, employing parallel programming to optimize throughput and enabling scalable automation pipelines for the entire engineering team.
- Automated hardware-in-the-loop (HIL) testing workflows with Python and MySQL-backed data persistence, integrating debugging and profiling to enable unattended execution and save 1000+ engineering hours annually.
- Designed and deployed a long-running, fault-tolerant test platform (ElectronJS + Django) to schedule and monitor 4000+ hour hardware validation runs with automated job recovery, email-based fault notifications, and task prioritization to ensure continuous part testing without manual intervention.
- Built CI/CD pipelines and containerized services on OpenShift, improving deployment reliability.
- Refactored existing codebases to implement isolated dev/prod MySQL environments, improving deployment reliability and data integrity.
- Initiated and built a secure ReactJS analytics platform to replace vulnerable Excel-based data sharing, safeguarding proprietary FET performance regressions via role-based access control.

**Jun 2025 - Dec 2025**

*Warwick, RI*

**Oracle**

*Software Engineer, Oracle Content Management*

- Developed and maintained large-scale backend services for content analytics and customer behavior tracking, integrating machine learning models to directly support lead generation and user engagement initiatives using Oracle CMS and ReactJS.
- Built scalable REST APIs backed by Postgres and distributed services to process high-volume content metadata and behavioral signals.
- Partnered with product and analytics teams to instrument performance metrics and optimize workflows, eliminating \$50K/year in outsourcing costs.
- Designed a custom adaptive-bitrate video player using VideoJS for Oracle Media Services, replacing third-party solutions and reducing recurring licensing costs.
- Implemented event-driven processing workflows using Kafka to support scalable content indexing and asynchronous job execution, applying deep learning techniques to enhance metadata extraction.

**Jul 2022 - Aug 2024**

*Remote, India*

## Projects

**MiniServe - A distributed ML Inference Pipeline** | <https://github.com/Kanchuking1/MiniServe>

**Feb 2026 - Present**

- Designed and developed a distributed asynchronous image classification inference service using **FastAPI** and **Redis Streams**, enabling scalable CPU-based model inference without blocking API threads.
- Architected a decoupled request pipeline where client uploads enqueue jobs (Redis Stream), and dedicated worker replicas asynchronously consume and process image classification, improving throughput and scalability.
- Deployed the full stack on an EC2 instance to deliver a publicly accessible demo link, demonstrating real-world operational readiness and cloud deployment proficiency.

**Faculty Teaching Assignment Software** | <https://github.com/tamu-edu-students/Faculty-Teaching-Assignment>

**Sep 2024 - Dec 2024**

- Built a linear-programming-based scheduling engine modelling faculty availability and academic constraints to generate conflict-free assignments, reducing scheduling time from ~4 weeks to ~1 week.
- Designed and deployed a Rails + PostgreSQL backend application on AWS, supporting 200+ course-faculty mappings per term with indexed schemas and optimized queries.
- Implemented RESTful endpoints and transactional database logic to ensure consistency and high availability.
- Led a team of 5 student developers to work together and managed as a communication point between the client and team allowing for smooth understanding of the requirements.

**C++-based Key-Value Storage Engine** | <https://github.com/Kanchuking1/key-value-server>

**Jan 2026 - Present**

- Designed and implemented a high-performance key-value storage engine with a non-blocking TCP server supporting hundreds of concurrent clients via poll(), optimized for sub-millisecond median latency.
- Implemented a dynamically resizing hash table with separate chaining and incremental rehashing, ensuring amortized O(1) lookups and stable latency during resize.

## Technical Skills

- **Languages:** C++17, Python, TypeScript/JavaScript, Bash, SQL, C/C++, Java
- **Systems & Networking:** Concurrency, TCP/IP, HTTP, MPI, Multithreading, Vectorization, debugging and profiling
- **Frameworks:** Django, FastAPI, React, Node.js, Ray, vLLM
- **Infrastructure, Cloud & Operations:** Docker, Kubernetes, OpenShift, AWS(EC2 & S3), CI/CD, Jenkins, Tableau, Alteryx, GitHub Actions, Redis, Redis Streams, Task Prioritization
- **Data & ML:** PostgreSQL, MongoDB, Redis, Kafka, PyTorch, CUDA, GEMM, Deep Learning, Machine Learning