Eshaan Bajpai

eshaangrad@gmail.com • linkedin.com/in/ishaanbajpai • github.com/Ishaan29 •

EDUCATION

Master of Engineering - Software Engineering 08/2024 - 05/2026

University of Maryland, College Park, United States

All Builders Welcome Grant Winner 2X - AWS re:Inforce, re:Invent 05/2025

Bachelor of Technology - Computer Science 08/2016 - 06/2020

SRM Institute of Science and Technology, Chennai, India

SKILLS

Languages: Python, Javascript, Java, Typescript, Kotlin, Go, Rust, C++, C#, Shell, SQL.

Engineering : Data Structures & Algorithms, Object-Oriented Design & Architecture (OOA), Distributed Systems, Database Systems, Operating Systems, System Design, Scalability, Reliability Engineering, Robust Systems, Observability & Monitoring, Test-Driven Development, Domain-Driven Architecture (DDA), Service-Oriented Architecture (SOA), Multi-Tenant Architecture, API-First Design, Microservices, ETL & Data Pipelines, Full Software Development Lifecycle, Production Operations, Security Best Practices, Automation.

Frameworks React, Angular, Spring Boot, Redux, Jest, Babel, JUnit, Hibernate, JDBC, Nodejs, Express, NestJs, Selenium, NPM, Maven, NVM, JUnit, TestNG, Selenium.

Database & Messaging Systems PostgreSQL (RDBMS), Delta Lake, GraphQL, Apache Spark, PGVector, Pinecone, MongoDB (NoSQL), Redis, Kafka, RabbitMQ, DynamoDB, S3.

Infrastructure & DevOps: AWS, GCP, Terraform, CloudFormation, EKS, ECS, ALB, WAF, AWS Bedrock, CI/CD, Kubernetes K8s, Helm, Docker, Opentelemetry, Sumologic, Grafana, Elastic Search, Azure Cloud, Jenkins, Splunk.

WORK EXPERIENCE

Senior Software Engineer | *Velotio, Inc* 03/2024 - 08/2024

- **Led full software lifecycle** of IoT telemetry platform processing **100K+ data points/sec** with sub-10ms latency and **99.9% reliability**.
- Architected **scalable data ingestion pipeline** using Java/Spring Boot and Kafka, achieving **30% throughput increase** for daily data.
- Deployed **highly available** AWS infrastructure (EKS, ALB) using Terraform, supporting **3x auto-scaling** with **zero downtime**.
- Implemented comprehensive testing strategy (unit, integration, load tests) achieving **85% code coverage** and **60% bug reduction**.
- Engineered TCP/IP and MQTT packet routing, enabling **150ms round-trip** device control latency.

Software Engineer Backend | *Gainsight, Inc* 02/2023 - 02/2024

- **Owned end-to-end development** of RAG-based analytics microservice processing **10TB data** across **400+ tenants** in beta with **99.95% uptime**.
- Implemented **reliable ETL workflows** with automatic retry mechanisms and data validation, achieving **2x performance improvement** through partition optimization.
- Built **scalable multi-tenant architecture** using Java, Python, and PGVector with **Delta Lake** for efficient data storage and retrieval.
- Improved system **observability and reliability** by implementing a custom OpenTelemetry SDK across 12 microservices, reducing MTTR by **50%**.
- Led Kubernetes migration for 5+ services, improving **deployment reliability by 60%** and eliminating environment drift, enabling development for **multiple data centers**.

Associate Software Engineer Backend | *Gainsight, Inc* 05/2021 - 02/2023

- Developed **high-reliability** notification engine serving **1M+ users** with 80ms p99 latency and **99.99% availability**.
- Built data pipeline integrations with Slack and Microsoft Graph APIs, processing **10K messages/sec** with **fault-tolerant design**.
- **Owned production stability**: diagnosed and resolved 200+ bugs, reducing recurring incidents by **80%** through root cause analysis.
- Implemented comprehensive test suites achieving **85% code coverage** across Node.js microservices using Jest and integration tests.

Software Requirements / QA Analyst Intern | *Center for Advanced Transportation Lab, UMD* 12/2024 - Current

- **Owned full testing lifecycle**, developed comprehensive unit tests (JUnit), integration tests, and automated E2E testing framework with (Cypress/Selenium), achieving **90% code coverage**.
- Built CI/CD pipeline with automated testing gates, reducing validation cycles by 30% and **preventing production defects**.
- **Delivered end-to-end features** from design through testing to production deployment, achieving **100% on-time delivery**.

PROJECTS

AI Desktop Assistant | *AI Memory That Actually Sticks* 05/2025 - Current

- **Leading development** of **multi-LLM orchestration** system with **intelligent request routing** across providers (OpenAI, Anthropic, Gemini), implementing **scalable caching layer** that reduced API costs by **40%**.
- Built **reliable context preservation** system enabling **seamless provider switching** without data loss, with **secure architecture** for screen sharing and semantic search using Pinecone.

VectorDB | *github.com/Ishaan29/vectorDB* 12/2024 - Current

- **Architected and built** distributed **vector database** from scratch supporting **100k+** vectors with **<100ms query latency**, implementing SIMD-optimized cosine similarity search achieving 5x performance improvement.
- Designed **fault-tolerant architecture** with consistent hashing, replication, and automatic failover ensuring **99.9% availability**.