Hyderabad, India
abhishekcp.singh67@gmail.com
linkedin.com/in/abhishekh-singh-636ab021
github.com/abhishekhsingh
abhishekhsingh.github.io

Abhishekh Singh

Principal Engineer

Summary

Principal Engineer with 12+ years of experience architecting and building production distributed systems at scale. Currently leading AI and search platform development across two product lines at Gainsight — designing data pipelines processing billions of records daily, AI-powered search with multi-provider LLM integration, and a production MCP (Model Context Protocol) server enabling AI assistants to query enterprise knowledge bases natively.

Experience

Apr 2025–Present **Principal Engineer**, *Gainsight*, Hyderabad

Architecting the Search & AI platform for Gainsight's Customer Communities product line. Leading the platform's evolution from traditional search to AI-native experiences.

- Built a production MCP (Model Context Protocol) server — **21** read-only tools exposing community knowledge to AI assistants (Claude, ChatGPT, Cursor) via SSE transport with JWT authentication and Redis-backed token management.
- Designed and shipped three LLM-powered features: an AI auto-reply system with cross-provider fallback (OpenAI + Gemini), an AI-powered answer engine with SSE streaming, and a search summary feature powered by Llama 3 70B.
- Architected a search analytics data lake on S3 (Parquet) + Athena, reducing upstream API costs by **90%**+ through query optimization and tag expression consolidation.
- Built a federated search system indexing **25M**+ records from 5 content sources with zero-downtime full reindex capability.

Python, FastMCP, Algolia, OpenAI, Gemini, Llama 3, AWS (ECS Fargate, S3, Athena, SQS), Redis, JWT, SSE

Aug 2019–Mar 2025 **Lead Software Engineer**, *Gainsight*, Hyderabad

Led architecture and development of core platform infrastructure across the Customer Success product line. Owned end-to-end delivery of production systems serving hundreds of enterprise tenants at scale.

- Architected the DAG-based data pipeline engine — processes **billions** of records daily across **600**+ tenants with a **98.7%** job success rate. Horizontally-scalable worker cluster with **35**+ pluggable SaaS connectors.
- Designed a four-tier event processing platform (Receiver → Router → Relayer → Broadcaster) with cross-region queue failover and self-healing delivery recovery.
- Built the analytics orchestration platform with dual-warehouse abstraction (PostgreSQL + Redshift), per-tenant dynamic schema generation, and ~**25** pluggable task builders.
- Delivered enterprise webhook system with IAM role assumption and request signing for secure cross-account delivery.

Java, Vert.x, Apache Drill, Zookeeper, Python, PostgreSQL, Redshift, Redis, AWS (SQS, S3, ECS), Docker

Jul 2017–Aug 2019 **Senior Software Engineer**, *Gainsight*, Hyderabad

Designed and built core platform services — the cross-cutting foundations that every other service depends on.

- Built a graph-based RBAC system on Neo4j — single graph traversal resolves access checks in milliseconds. Permission SDK consumed by **7**+ platform services for row-level data enforcement.
- Designed the centralized configuration service — runtime config with encrypted secrets, version history, and periodic refresh. Config client SDK embedded across all platform Java services.
- Built the reactive audit service — fire-and-forget writes decoupling caller latency from persistence, with lazy per-tenant table creation and separate read/write connection pools.
- Developed **60**+ custom SQL functions for tenant-specific data transformations in the pipeline engine.

Java, Vert.x, Spring Boot, Neo4j, Redis, PostgreSQL, Apache Drill

Jul 2016–Jun 2017 **Software Engineer**, *Gainsight*, Hyderabad

Joined as the company's data platform was being built from the ground up. Contributed to foundational infrastructure that still powers the platform today.

- Contributed to the initial architecture of the graph-based permission system and the distributed configuration service.
- Built data transformation components and connector integrations for the pipeline engine.
- Developed shared platform libraries used across multiple services.

Java, Vert.x, Neo4j, PostgreSQL, Redis

Dec 2015–Jun 2016 **Software Engineer**, *Talentica Software*, Pune

Engineered the Intelligence Machine — a big data recommendation engine processing terabytes of social media data daily for content personalization. Built performance benchmarking infrastructure for distributed systems, including load simulation and bottleneck analysis across multi-node deployments.

Java, Big Data, Performance Engineering, Distributed Systems

Aug 2014–Dec 2015 **Software Analyst**, *Nihilent*, Pune

Delivered 5 projects for MultiChoice Africa (Naspers subsidiary) — one of Africa's largest media platforms serving millions of subscribers. Built automation frameworks using Selenium and UFT for end-to-end regression coverage. Interfaced directly with onshore client teams in South Africa to ensure delivery quality.

Selenium, UFT, JMeter, Client Delivery

Technical Skills

Languages	Java, Python
Frameworks	Vert.x, Spring Boot, FastAPI, FastMCP
AI / ML	OpenAI, Google Gemini, Llama 3, Algolia AskAI, MCP Protocol, SSE Streaming
Cloud & Infra	AWS (ECS Fargate, Lambda, S3, SQS, SNS, Athena, Redshift, DynamoDB), Docker
Data & Search	PostgreSQL, Redis, Neo4j, Algolia, Apache Drill, MongoDB, MySQL, Zookeeper
DevOps	GitHub Actions, Bitbucket Pipelines, Maven, SonarCloud, Git

Education

2010–2014 **B.Tech. Computer Science & Engineering**, *BBD National Institute of Technology & Management*, Lucknow, 75.16%

Honors & Achievements

- Won **nine** spot awards at Gainsight across 5 years for consistent delivery and engineering excellence.
- Vice-Captain, Nihilentia Annual Technology Fest (2014) — led event coordination across 200+ participants.
- First Prize in Cyber Attack Finale at UTKARSH-13, BBDNITM Lucknow.
- Third position in UP Board High School Examinations, Jaunpur Block (2007).