

MURALIKRISHNA DUDAKA

Senior Data & AI Engineer

Cloud Lakehouse | Streaming Analytics | Applied AI

✉ 2dmurali@gmail.com

☎ +1 669 292 8805

📍 Santa Clara, CA



Professional Summary

- Senior Data & AI Engineer with over 15 years of experience designing, deploying, and optimizing hyperscale cloud environments.
- AWS & Databricks Certified expert specializing in resilient cloud-native lakehouse platforms using Apache Spark, Flink, and Kafka on Kubernetes.
- Proven track record of processing petabytes of streaming and batch data with sub-second latency while driving substantial reductions in cloud operating expenditures.
- Experienced in integrating modern AI tools into enterprise infrastructure, including agentic frameworks and the Model Context Protocol (MCP) for complex data operations.

Technical Skills

- **Cloud & Infrastructure:** AWS (EKS, S3, RDS, Lambda), Kubernetes, Docker
- **Data Engineering & Streaming:** Databricks, Apache Spark, Apache Flink, Apache Kafka, Apache Iceberg, Apache Pinot
- **AI Integrations:** Large Language Models (LLMs), Model Context Protocol (MCP), Applied Machine Learning
- **Software Engineering:** Java, Python, Spring Boot

Professional Experience

Cisco Systems | Software Engineer (Senior/Staff Data Engineer) | Sep 2024 - Present

- Engineered a petabyte-scale streaming analytics framework using Apache Flink and Apache Pinot to modernize core enterprise contact-center reporting.
- Guaranteed sub-second query latency for more than 10,000 global customers.
- Reduced continuous compute costs by \$150k-\$200k monthly and cut pipeline processing time by 80% through optimized streaming ingestion patterns.
- Architected an agentic operational interface for Apache Pinot using LLMs and MCP, enabling cross-functional teams to query real-time streaming data through natural language.
- Developed a globally distributed Data Federation architecture across multiple regions, integrating batch and streaming pipelines into Apache Iceberg with GDPR-compliant PII masking.
- Leveraged the DeltaIO framework to build a customized Bring Your Own Compute (BYOC) infrastructure, enabling more than 800 enterprise customers to execute isolated analytical workloads.
- Led architectural review boards to establish enterprise-wide standards for streaming data platforms and AI integrations.

Infosys Limited (Client: Apple Inc.) | Senior Cloud Big Data Engineer | Jan 2015 - Sep 2024

- Engineered and deployed a homegrown multi-cloud data lake ecosystem from the ground up to support Apple's petabyte-scale operational demands.
- Spearheaded the migration of massive data workloads from the on-premise Apple Private Cloud to an auto-scaling Amazon EKS infrastructure.

Professional Experience (continued)

Infosys Limited (Client: Apple Inc.) | Senior Cloud Big Data Engineer (continued)

- Integrated Apache Spark and Apache Iceberg on Amazon S3 to analyze petabytes of data with strict ACID compliance.
- Designed an active-active, cross-region Disaster Recovery framework in AWS.
- Created a custom S3 data replication engine that bypassed native replication features, reducing disaster-recovery infrastructure costs by millions of dollars annually.
- Achieved an 80% reduction in cloud compute costs and a 4x improvement in query execution performance through rigorous Apache Iceberg table rewrite procedures.
- Mentored and upskilled mid-level and junior engineers in cloud-native data engineering and distributed systems best practices.

Tata Consultancy Services (TCS) | Senior System Engineer | Mar 2011 - Dec 2014

- Led the real-time, secure integration of India's Aadhaar biometric system with legacy core-banking infrastructure for Andhra Bank.
- Designed and implemented RESTful web services using Spring Boot for banking and government-sector applications.
- Architected and developed an operational portal for NMPT to automate ground-level cargo handling and secure regulatory communications.

Independent Projects & Open Source

- **Automated AI Trading Platform:** Built and deployed an end-to-end cloud-native trading platform featuring a multi-model machine-learning ensemble for market data and institutional price-action analysis.
- **Apache Pinot Contribution (PR #17221):** Contributed code optimizations to the official open-source codebase, resolving JVM resource-utilization issues during on-demand administrative processes.
- **OpenClaude Contribution (PR #991):** Contributed successfully merged updates to the open-source repository.

Certifications & Education

- AWS Certified Solutions Architect - Professional
- Databricks Certified Data Engineer Professional
- Machine Learning Specialization (Completed: April 2026)
- Master of Computer Applications (MCA), National Institute of Technology (NIT), Raipur, India (2010)
- Bachelor of Science in Computer Science, Acharya Nagarjuna University, India (2007)

Publications

- A Privacy-Preserving Big Data Frameworks for Ensuring Secure and Efficient Network Communication in Distributed Systems. *IEEE*, 2026.
- A Robust Framework for Big Data Security in Distributed Network Environments. *IEEE*, 2026.
- Dual-Write Reconciliation Between Real-Time Serving Stores and Open Lakehouse Tables: A Comparative Study of Three Strategies. *AES*, 2026.
- Tool-Surface Granularity in Model Context Protocol Servers for Analytical Engines: Patterns, Trade-Offs, and Token Economics. *International Journal of Intelligent Systems and Applications in Engineering (IJISAE)*, 2026.