RISHABH GUPTA

Solution Architecture | Data & AI Platforms | Cloud & Systems Integration
Gurugram, India | em.rishabh@gmail.com | +91 88262 89159 | linkedin.com/in/emrishabh

Professional Summary

Technology and Data Platform Engineer with 10+ years of experience designing, integrating, and operating enterprise applications, cloud infrastructure, data platforms, AI-enabled solutions, and production services. Experienced in translating business and operational requirements into scalable solution designs, integration patterns, data models, validation controls, APIs, and support models.

At adidas, serve as a technical owner for the Databricks layer of the M3 platform, supporting architecture, Spark transformation and optimization, data quality, Kafka integration, cloud automation, observability, and production reliability across Spark and Delta workloads. One of the designers and developers of a deployed agentic AI assistant for support, integrating RAG and enterprise MCP tools for incident investigation, root-cause analysis, and post-mortem support. Provide technical leadership to a devops and SRE engineer team through design reviews, code reviews, prioritization, troubleshooting, and mentoring.

Core Competencies

Solution Architecture:
Solution design, systems integration, interface design, cloud/data architecture, option and trade-off analysis, non-functional requirements, technical specifications, design reviews

Data & Lakehouse:
Databricks, Apache Spark, PySpark, SQL, Delta Lake, Unity Catalog, DLT/Lakeflow, data modelling, ETL/ELT, data governance, data quality

AI & Analytics:
Agentic AI, tool calling, routing logic, RAG, MCP integration, incident intelligence, Databricks Genie POC, LightGBM, feature engineering, model validation and explainability

Integration & APIs:
Kafka, Kafka Connect, Schema Registry, REST APIs, Python, Flask, FastAPI, JDBC, Jira, Confluence, ServiceNow, Kubernetes and Instana integrations

Cloud & Databases:
AWS, Azure, S3, Lambda, IAM, CloudWatch, Terraform, Kubernetes, Aurora MySQL, SQL Server, DocumentDB, Oracle source investigation

Leadership & Delivery:
Requirements clarification, Agile delivery, stakeholder management, prioritization, code/design reviews, mentoring, incident response, RCA, technical documentation

Professional Experience

adidas | July 2022 -- Present

Software Engineer II ? Data Platform & Site Reliability | Gurugram, India

? Serve as technical owner for the Databricks layer of the M3 enterprise platform across 20+ Spark and Delta workloads.
? Translate business, operational, platform, and data requirements into solution designs, data models, integration flows, validation logic, reusable pipelines, APIs, and implementation plans.
? Design Lakehouse structures using Delta Lake and Unity Catalog and review Databricks Jobs, Workflows, and DLT/Lakeflow pipelines for ingestion, transformation, validation, reconciliation, and delivery.
? Optimized a production pipeline processing approximately 600M records, reducing runtime from approximately 10 hours to 4 hours through partitioning, join optimization, skew handling, caching, and data-layout improvements.
? Co-designed and developed a deployed agentic AI assistant for support using RAG retrieval, enterprise MCP tools and multi-step investigation workflows.

? Integrated Jira, Confluence, ServiceNow, Kubernetes, and Instana through MCP tools to enrich incident investigation, root-cause analysis, and post-mortem analysis.
? Built integration capabilities using Kafka Connect, Schema Registry, topic provisioning, ACL management, consumer-lag monitoring, REST APIs, and JDBC.
? Automated AWS infrastructure using Terraform across S3, Lambda, IAM, and CloudWatch and built operational dashboards with OpenSearch, Grafana, Instana, and CloudWatch.
? Provided technical leadership for an 8-engineer team through solution and code reviews, prioritization, troubleshooting, mentoring, and engineering standards.

Technology: Databricks, PySpark, Spark SQL, Delta Lake, Unity Catalog, Python, SQL, LightGBM, RAG, MCP, Kafka, AWS, Terraform, Kuberr

Optum (UnitedHealth Group) | January 2020 -- July 2022

Senior Software Engineer | India

? Collaborated with business, application, infrastructure, and deployment teams to modernize enterprise application and data platforms.
? Translated functional and operational requirements into technical implementation plans, ETL workflows, deployment activities, and integration changes.
? Led Windows Server and Azure infrastructure migration initiatives, coordinating application, infrastructure, security, deployment, and release dependencies.
? Built deployment automation and designed NGINX reverse-proxy architecture, supporting integration, release reliability, operational readiness, and production troubleshooting.

Optum (UnitedHealth Group) | January 2018 -- December 2019

Software Engineer | India

? Built SQL Server SSIS ETL pipelines and SSRS reporting solutions for enterprise ingestion, transformation, validation, and reporting workloads.
? Designed and optimized SQL queries, stored procedures, indexes, and reporting structures to improve processing and reporting performance.
? Implemented IBM Tivoli Workload Scheduler workflows with dependency management, monitoring, failure handling, and operational support.
? Investigated production data-processing failures, validated root causes, and implemented corrective actions with application and business teams.

UnitedHealth Group | July 2016 -- January 2018

Application Engineer | India

? Supported enterprise applications through backend development, deployments, integration support, defect resolution, SQL analysis, and production troubleshooting.

Selected Architecture, AI & Analytics Projects

Agentic AI Assistant for L1 Support

adidas ? Deployed incident-investigation and operational intelligence assistant

- ? Co-designed and developed an agentic assistant that routes support questions across RAG retrieval, enterprise tools, and multi-step investigation workflows.
- ? Integrated Jira, Confluence, ServiceNow, Kubernetes, and Instana MCP tools to gather incident, documentation, workload, and observability evidence for enriched analysis.
- ? Built routing and tool-calling logic for historical-incident retrieval, issue enrichment, root-cause investigation, and post-mortem analysis while retaining source evidence and operational context.
- ? Supported deployment and ongoing development, including prompt and workflow refinement, integration improvements, and support-use-case expansion.

Capabilities: Agentic AI, RAG, tool calling, MCP integration, incident analysis, RCA, post-mortem support

M3 Data Platform & Supply-Chain Integration

adidas ? Enterprise Databricks, Lakehouse, integration, and operational data platform

- ? Defined and supported Databricks architecture for ingestion, transformation, validation, reconciliation, and downstream delivery, integrated with Kafka, AWS, SQL Server, and operational tools.
- ? Established governed Lakehouse structures and data-quality controls and applied Spark performance engineering to improve scalability, reliability, runtime, and production operability.
- ? Built M3-to-SAP AFS recovery automation across BWCL/RMA, TIBCO, SAP AFS acknowledgement flows, Lakehouse, and SQL Server evidence sources.

Technologies: Databricks, Spark, PySpark, Delta Lake, Unity Catalog, SQL, Kafka, AWS, Terraform

Range-Selection Recommendation Model & Databricks Genie POC

adidas ? Supply-chain decision support and conversational analytics

- ? Built a LightGBM model scoring approximately 2.6M article-market-season combinations across 11 markets.
- ? Achieved validation results of approximately 0.87 AUC and 0.96 PR-AUC and designed confidence bands, explainable recommendation tiers, and reason codes for business review.
- ? Reconciled model recommendations with live merchandising decisions to identify aligned cases, pending decisions, and high-priority disagreements.
- ? Developed a Databricks Genie POC for natural-language exploration of ranging and sizing data and recommendation outputs.

Technologies: Databricks, PySpark, SQL, LightGBM, Delta Lake, Apache Spark, Databricks Genie

EDC Balance

adidas ? Database-backed application and scheduled operational workflows

- ? Designed the application architecture, Python backend services, Aurora MySQL schemas, and scheduled reconciliation workflows; integrated application, database, and AWS services.

Technologies: Python, SQL, Aurora MySQL, AWS

Recognition & Education

UnitedHealth Group Bravo Awards | 2019 -- 2022

Awards recognizing cloud migration, claims-workflow optimization, technical delivery, and cross-functional contributions.

B.Tech., Computer Science & Engineering | 2012 -- 2016

Dr. A.P.J. Abdul Kalam Technical University