



DE Delivery Experience: 15 Years

Headquarters: **Texas, USA**

Overall Employee Size: **1000+ FTEs**

DE Delivery Team Size: **100+ FTEs**



[Website](#)



[LinkedIn](#)

Data Engineering Services Overview

DE Services

Qubika delivers end-to-end data engineering services anchored in a lakehouse-first, AI-enablement philosophy. Across strategy and architecture, ingestion, platform engineering, governance, analytics/AI/GenAI enablement, DataOps/MLOps, and managed services, the firm operates a Medallion architecture-driven delivery model predominantly on Databricks, Snowflake, AWS, Azure, and GCP. Qubika's primary architectural pattern follows Bronze-Silver-Gold layering with idempotent pipeline design, reusable ingestion frameworks, multi-environment CI/CD, and cost-aware compute governance. Both batch and streaming workloads are supported using Spark Structured Streaming, Kafka, Databricks Auto Loader, and Airflow/Databricks Workflows, across cloud-native, on-premise, and hybrid deployments. The firm's workload mix skews heavily toward core data engineering and data science (72%), complemented by GenAI/agentic AI enablement (19%) and BI/analytics (9%). Majority of engagements incorporate AI-assisted delivery, with production workloads operating on a supervised hybrid model. In FY2025, Qubika delivered 51 data engineering engagements. The majority of client engagements are delivered via T&M, along with fixed-price and outcome-based pricing models.

DE accelerators, and toolkits

- **QBricks** - Enterprise agentic AI platform built on LangGraph and Databricks, centralizing the full agent lifecycle (build, evaluate, deploy, monitor) within a governed, lakehouse-native environment with pre-built templates and a cost/performance observability dashboard.
- **Governance accelerators** - *Unity Catalog Setup Accelerator* and *Governance Migration Agent* automate Databricks governance setup and metadata migration from Snowflake and Redshift, cutting effort by up to 90%.
- **Data-to-AI accelerators** - *MultiFormat Parsing Agent*, *Unstructured Knowledge Assistant*, and *GraphRAG* convert unstructured content into governed, AI-ready Delta tables, knowledge graphs, and semantic search/RAG workflows.
- **Business and FinOps accelerators** - *Finance Analyst AI Agent* reduces insight generation time and reporting effort, while the *Databricks FinOps Accelerator* (in development) provides unified cost visibility, forecasting, and optimization recommendations.

Markets served

North America (~90%), Latin America (~5%), EMEA (~5%)

Industries served

Financial Services, Media & Entertainment, Technology, Healthcare and Education

Key partnerships

AWS (Advanced Tier), Databricks (Elite/Gold Partner), Snowflake (Select Partner)



Key Highlights

Strengths

- **Fastest-growing Databricks Gold-tier engineering partner in the Americas:** With 200+ Databricks-certified professionals and Gold/Elite partner status, Qubika operates as a Databricks co-sell, direct-client, and Databricks Delivery Partner Program (DPP) simultaneously - a structurally rare tripartite GTM position giving it deep ecosystem access and joint pipeline visibility unavailable to most competitors.
- **Proprietary accelerator portfolio delivering 90% reduction in platform setup and governance migration effort:** The Unity Catalog Setup Accelerator and Governance Migration Agent Accelerator directly compress the highest-friction phases of modern lakehouse adoption, enabling faster client time-to-value than manual implementation approaches and supporting 25+ data migration and platform modernization projects in FY2025, alongside Qubika's strong YoY DE growth trajectory.
- **AI-native, senior-led delivery model generating measurable productivity advantage:** 84% of the delivery team is mid-career or above (the "diamond AI cell" model), combined with AI-assisted delivery across all engagements, generating a reported 40%+ productivity uplift.
- **Strong production reliability and delivery accountability:** 95% client retention, very high SLA/delivery target adherence across active engagements, and MTTR of 30 minutes to 2 hours in production environments, supported by automated monitoring, idempotent pipeline design, and structured on-call rotations.
- **Dual-certified security posture with NIST AI RMF compliance:** SOC 2 Type II, ISO/IEC 27001:2022, and alignment with the NIST AI Risk Management Framework, positioning Qubika well for secure, regulated data and AI environments.
- **Deep agentic AI and GenAI production capability embedded in DE delivery:** Beyond traditional pipeline work, Qubika operationalizes Vector Databases, Knowledge Graphs, RAG pipelines, and multi-agent architectures as core DE infrastructure enabling clients to move from data platform modernization to GenAI production in a single engagement arc.
- **Concentrated vertical expertise with domain-specific IP:** Significant production depth in financial services (payment processing, credit scoring, compliance automation, SOX reconciliation) and media/entertainment underpins reusable solution frameworks and accelerates domain-specific client value delivery.

Case Study Highlights

Modernizing a Financial Data Backbone for a Global Payments Firm

A global payments platform faced fragmented financial data systems, causing a 5-day bank close cycle, ~10,000 monthly analyst hours of manual effort, and compliance risk from 1,500+ controls. Qubika built a modern financial data platform using Redshift, EMR, Spark, Airflow, and Apache Iceberg, automating volatility classification, reconciliation, and audit workflows through rules-based logic and LLM-assisted SQL agents. The engagement delivered a 40% reduction in bank close cycle, 90% reduction in manual reconciliation effort, 85% fewer compliance errors, and ~\$60K+ in monthly savings while improving auditability and real-time data access.

Upgrading a Mission-Critical Lakehouse Platform for a Global Airline

A global airline faced operational risk as a legacy Databricks runtime approached end-of-life across production pipelines tightly coupled with ML workloads. Qubika executed a phased platform upgrade across DEV, UAT, and PRD environments, with full dependency standardization, compatibility resolution, and rollback safeguards at each stage to ensure continuity. The engagement eliminated end-of-life infrastructure risk with zero production disruption, standardized configurations across all environments, and delivered a more resilient, governance-ready data foundation to support ongoing AI and analytics innovation at scale.