

AIM Research
Penetration Maturity (PeMa) Quadrant

Top Databricks Services Partners 2026

AIM Research's annual initiative recognizes leading Databricks services partners that help enterprises build, migrate, and operationalize data and AI workloads on the Databricks Lakehouse Platform, while providing buyers with a trusted benchmark for partner evaluation.

© 2026 AIM Media House LLC and/or its affiliates. All rights reserved. For more information, email info@aimmediahouse.com or visit aimresearch.co.

Licensed to: Qubika | License ID: TDB/PeMa/AIM/Qubika-26





Qubika

Founded year: **2005**

HQ: **Austin, Texas**

Client Base: **US, Canada, Latin America, and Europe**

Partner Type: **Consulting & SI**

Tier: **Gold**

Databricks Certifications: **250+**



[Website](#)



[LinkedIn](#)

Databricks Practice Overview

Lifecycle Coverage	Qubika provides end-to-end services on the Databricks Lakehouse Platform, supporting clients from data foundation to advanced AI and enterprise-grade agents. Foundational build and modernization work makes up the bulk of Qubika's Databricks engagements (roughly 70% across lakehouse, engineering, migration, and BI), complemented by steady delivery across advanced analytics, Generative and Agentic AI, governance, and platform operations.
Top Verticals	Gaming & Sports, Media & Entertainment, Banking, Financial Services & Insurance, and Manufacturing & Industrials
IP & Accelerators	QBricks - A Built on Databricks platform that standardizes the design, deployment, and governance of AI agents at scale, with orchestration, evaluation, observability, and cost tracking built in. Finance Analyst AI Agent - Delivers fast, accurate financial insights to executives from organizational data that would traditionally take a data analyst hours or days to prepare. Unity Catalog Setup Accelerator - Automates governance setup using LLM parsing and metadata mapping to deploy secure, standardized Unity Catalog environments. Governance Migration Agent Accelerator - Automates migration of governance metadata from Snowflake and Redshift into Unity Catalog, extracting, validating, and translating objects into Databricks-ready definitions. Unstructured Knowledge Assistant - Enables semantic search and AI Q&A over document repositories within Databricks, automating ingestion, embedding, and governance with full traceability. MultiFormat Parsing Agent Accelerator - Ingests, parses, and structures multi-format content including PDF, HTML, image, audio, and video into AI-ready Delta tables. GraphRAG Accelerator - Transforms unstructured and semi-structured data into governed knowledge graphs and retrieval pipelines within Databricks. MLOps Accelerator - Automates the model lifecycle through CI/CD, model tracking, and periodic health reviews. Qubika AI Dev Kit - An AI coding assistant that builds Databricks pipelines with catalog naming, medallion architecture, Unity Catalog governance, guardrails, and testing applied by default. Reported to cut time-to-value by 40 to 70%.



Key Highlights

Partnership standing: Databricks awarded Qubika its Specialization for Communications, Media, Entertainment & Gaming in July 2026, a credential granted against evidenced sector engagements and sitting above Gold Tier partner status. The firm reached Gold from Select in roughly two and a half years, tripling joint Databricks customers along the way, with a Data & AI Studio of 300-plus engineers behind it, 250 of them Databricks-certified and three holding Champion status, the platform's invite-only tier for senior technical experts.

Migration track record: More than ten financial services clients have been migrated to Databricks to date, with delivered work spanning Redshift, Dremio, Snowflake, Microsoft Fabric, BigQuery, and on-premise sources. The largest is a gaming and betting platform serving over 18 million customers, where the Gold layer is live and Bronze and Silver are in flight, alongside terabyte-a-day Kafka streaming into Delta and regulatory reporting rebuilt on the Lakehouse.

Recovery of stalled programs: At one lending client, two earlier migrations had collapsed under a prior vendor whose 50-odd person team was buried in constant problems. Qubika rebuilt the platform with about 10 senior engineers, who then went on to lead the client's internal teams. Operations stabilized, and the account has since expanded into credit scoring, customer journey analytics, and dispute automation.

Unity Catalog and governance: Net-new provisioning was cut by 90% in one lakehouse build, governance metadata was migrated off Snowflake for a mortgage lender with 80% less manual effort and onboarding compressed from months to weeks, and governance was unified across every operating unit of a community bank at fully aligned environments and 70% less manual oversight. A separate engagement covered multi-workspace governance across development, UAT, and production environments integrated with externally hosted models.

Accelerators and IP: Client code and IP stay portable at exit, with accelerators shipping through Databricks Asset Bundles so everything is packaged and version-controlled inside the platform, and agents built on QBricks remaining runnable in any comparable LangGraph environment without it. Nothing proprietary has to be unwound to leave, and the work stays legible to an internal team or a successor vendor.

AI and machine learning in production: Agentic work includes a retrieval system for a university endowment manager that Databricks called innovative at a Data + AI Summit, dispute automation running governed ML at 99.9% adjudication consistency, and an advertising-order solution for a broadcast media company built on LangGraph and Databricks Model Serving. On the ML side, a cross-sell model for a financial-services marketing platform outperforms random recommendations by up to 740%, and a lending client runs 15-plus models in production with credit decision latency down from days to seconds. Delivery covers the full lifecycle through Feature Store, Model Registry, Model Serving, and MLflow.

Delivery bench and staffing model: Delivery teams are staffed senior, with engineers on the lending recovery each carrying four or more years of Databricks experience. Intake runs through a dedicated training pipeline that drew 1,000-plus applications from Latin American universities, admitted 60 of 100 pre-selected candidates to intensive training, and retained 95%. Over 99% of staff sit in the Americas, with engineering concentrated in Latin America and same-day overlap with North American clients.

Run cost and platform operations: Cost optimization is offered as a consulting service, with outcomes including a 40% operational cost reduction through Databricks SQL and a 60% cut in data initiative cost at one lending client. Operations extend to runtime lifecycle work, including a phased runtime upgrade at a flag-carrier airline with rollback safeguards and weekly executive risk tracking.