

Data Cosmos_Genie AI

Conversational Analytics Chatbot on Databricks MLOps

Part of the Coforge Data Cosmos™ Platform | Quasar™ Portfolio

Coforge Limited | Data HBU

2026

Solution Overview

Data Cosmos_Genie AI is an intelligent, web-based analytics chatbot built on the Cosmos-ML-Flow Airlines MLOps platform powered by Databricks. It enables business users to explore airline booking fraud detection, flight delay prediction, and ML model health data through a simple conversational interface — no SQL or data engineering skills required. Part of Coforge Data Cosmos™, the innovation backbone combining platforms, agentic accelerators, and services, it democratizes access to advanced analytics and ML insights across the enterprise.

In modern enterprises, business analysts and operations teams depend on data engineering teams to write SQL, build reports, and surface insights — creating bottlenecks, ticket queues, and delayed decisions. Genie AI removes this dependency by translating plain-English questions into optimized SQL through Databricks Genie, backed by production ML models that score booking fraud and flight delays in real time and batch. Whether identifying fraudulent transactions with 90%+ detection rates, predicting delays hours before departure, or monitoring model health for drift, Genie AI delivers instant, SQL-backed answers with interactive visualizations and full query transparency — all through a secure conversational interface.

Key Features

Feature	Description
Natural Language Querying (NLP-to-SQL)	Users ask questions in plain English; Databricks Genie AI translates them into optimized SQL queries and returns results instantly — eliminating the need for SQL expertise.
13 Pre-Built Quick Queries	One-click sidebar buttons for the most common analytics questions covering fraud summary, delay analysis, booking volume, weather impact, model metrics, and more.
Interactive Data Visualization	Responses render as formatted HTML tables with sticky headers and Chart.js visualizations (bar, line, doughnut charts) for clear, actionable data views.
SQL Transparency & Dual-Source Badging	Every response includes an expandable SQL toggle to inspect the exact query executed, and is tagged with a 'Genie AI (NLP)' or 'Quick Query' badge so users know how each answer was generated.
Real-Time & Batch ML Scoring	Underlying ML models (LightGBM champion, XGBoost challenger) score booking fraud and flight delays on automated schedules with automated quality gates for model promotion.
Drift Detection & Model Monitoring	Continuous tracking of model accuracy, precision, recall, F1, and feature drift ensures models remain reliable in production, with secure Flask session-based authentication controlling access.

Six-Step Interaction Flow

Step	Role	What It Does
User Authentication	Secure Access	Users log in via a secure web interface. Flask manages sessions with a secret key, ensuring only authorized users access the platform.
Chat Interaction	Query Entry	Users either click a Quick Query button in the sidebar or type a free-form question in the chat input box.
Intent Matching	Keyword Routing	The backend scans the message against 13 pre-defined query templates. If keyword coverage exceeds 30% (or the message is ≤ 5 words), the matching Quick Query runs directly against the SQL Warehouse.
Genie AI NLP Fallback	NLP-to-SQL Generation	If no Quick Query matches, the message is routed to the Databricks Genie Space REST API, which generates SQL from the natural language question, executes it, and returns the result.
Response Rendering	Visualization & Badging	The result is formatted as an interactive HTML table with optional Chart.js visualizations and an expandable SQL preview, delivered with a source badge indicating how the answer was produced.
Continuous ML Pipeline	Backend Automation	Automated Databricks workflows run model training (daily), batch inference (every 4–6 hours), and monitoring (daily) to keep predictions fresh and accurate across fraud and delay use cases.

How Genie AI Works

1. **Authenticate** — Users log in via a secure web interface with Flask session management and a secret key
2. **Interact** — Users click a Quick Query sidebar button or type a free-form question in the chat input box
3. **Match Intent** — Backend scans the message against 13 query templates; $>30\%$ keyword coverage runs the Quick Query
4. **Execute Quick Query** — Matched Quick Queries run directly against the Databricks SQL Warehouse for instant results
5. **NLP Fallback** — Unmatched questions route to the Databricks Genie Space REST API for NLP-to-SQL generation
6. **Generate SQL** — Genie AI starts a conversation, generates SQL from natural language, and executes it
7. **Render Response** — Results format as interactive HTML tables with Chart.js visualizations and SQL preview
8. **Badge Source** — Each response is tagged 'Genie AI (NLP)' or 'Quick Query' for full transparency
9. **Score ML Models** — Backend LightGBM/XGBoost models score fraud and delays in real-time and batch
10. **Monitor & Retrain** — Automated training (daily), inference (4–6 hrs), and monitoring keep predictions accurate

Industry Trends and where it can be leveraged

Banking & Financial Services

Use Case	Example
Conversational Fraud Analytics	A global bank uses Genie AI to let fraud and risk teams explore transaction fraud patterns in plain English. NLP-to-SQL surfaces risky transactions and payment-method fraud analysis instantly, with ML models scoring transactions at 90%+ detection rates versus 60–70% for rule-based systems.
Self-Service Risk Exploration	A retail bank empowers business analysts to ask ad-hoc questions about transactions and model outputs without raising tickets — with SQL transparency building trust in every answer.

Insurance

Use Case	Example
Claims Pattern Exploration	A multi-line insurer uses Genie AI to let claims and analytics teams explore claims patterns and anomalies conversationally. Quick Queries provide one-click access to KPIs while NLP handles ad-hoc questions without SQL.
Model Health Transparency	An insurer's actuarial teams monitor model accuracy, precision, and drift through the chatbot, ensuring predictive models remain reliable before degradation impacts underwriting decisions.

Travel, Transportation & Hospitality

Use Case	Example
Booking Fraud & Delay Prediction	A global airline uses Genie AI's native use cases — booking fraud detection (14+ behavioral features) and flight delay prediction (weather, carrier, aircraft age, route history) — enabling proactive passenger rebooking and crew reassignment hours before departure.
Executive & Operations Dashboarding	A hotel and travel group uses Quick Query buttons for one-click KPIs — daily booking volume and revenue, top risky routes, delay rates by carrier, and weather impact on operations.

Healthcare & Life Sciences

Use Case	Example
Conversational Operational Analytics	A healthcare provider adapts Genie AI to let operations analysts explore patient flow, capacity, and operational metrics in natural language — removing dependency on data teams for routine reporting.
Model Monitoring for Clinical ML	A pharma company uses Genie AI's drift detection and model health monitoring to track clinical and research model performance, ensuring retraining before accuracy degrades.

Cross-Industry Use Cases

Use Case	Description
Ad-Hoc Data Exploration	Business analysts ask any question about data in natural language and get instant SQL-backed answers — no tickets, no waiting for reports.
Self-Service BI & Dashboarding	Quick Query buttons provide one-click access to KPIs and metrics, democratizing dashboarding for executives and operations teams.
Fraud Detection Analytics	Explore ML-scored fraud predictions conversationally, achieving 90%+ detection versus 60–70% with traditional rule-based systems.

Use Case	Description
Predictive Operations Planning	Surface ML delay/risk predictions in plain English to enable proactive operational and resource planning hours in advance.
ML Model Health Monitoring	Provide real-time visibility into model accuracy, precision, recall, F1, and feature drift to trigger retraining before degradation.

Target Audience

- Business Analysts & Operations Teams
- Data Science & ML Engineering Teams
- Fraud, Risk & Compliance Teams
- Executive & BI Leaders
- Data Engineering & Platform Teams

Azure Compatibility

Data Cosmos_Genie AI is built on Azure Databricks (Genie Space, SQL Warehouse, Unity Catalog) as its core MLOps and analytics platform, with Azure OpenAI Service augmenting natural language understanding. It is Dockerized with Gunicorn support and deploys on Azure Kubernetes Service (AKS) or any container platform with minimal configuration. Azure Data Lake Storage provides the Delta Lake data layer, Azure Monitor delivers observability, and Microsoft Entra ID can extend the Flask session-based authentication for enterprise SSO. It is part of the Coforge Data Cosmos Quasar portfolio for conversational AI and advanced analytics on databricks, fully designed to leverage Azure-native services for scalable, secure, and governed conversational analytics.