

Deloitte.

 databricks



TRUST AT SPEED

Deliver traceable, regulatory-ready answers with clear governance and data lineage using a curated Databricks Genie space



Executive Summary

Pharmaceutical manufacturing inspection readiness often breaks down not because data is unavailable, but because it is distributed, inconsistently defined and time-consuming to investigate. This means site teams must bridge the gap between dashboard key performance indicators (KPIs) and the next layer of inspection questions: What changed and why? Which lots were impacted? Who approved? What evidence supports the conclusion?

Deloitte's perspective is that a governed conversational layer can close that gap. Using Databricks Genie, manufacturing and quality teams can ask questions in plain language about governed, permissioned data and get audit-ready answers backed by approved sources and reusable structured query language (SQL). This document outlines how Deloitte helps design and scale Genie spaces with the right performance, security and change control so investigations can move faster and inspection narratives stay consistent—without weakening governance.

This paper is designed for pharmaceutical manufacturing, quality, and data leaders preparing for (or supporting) regulatory inspections. It focuses on how a governed conversational experience can help teams designing a curated Databricks Genie space to defensible, inspection-ready narratives. This governed conversational experience is not a replacement for validated quality systems, formal investigations, or established review/approval workflows; instead, it is an acceleration layer for finding, connecting, and explaining evidence. Examples in this paper are illustrative and should be tailored to each organization's quality management system and data controls.

Introduction:
The “conversation” starts here

For decades, organizations have interacted with data through rows, columns and carefully structured queries. That model has served us well, but it also shapes who can ask questions, how quickly they can iterate and how easily insights guide day-to-day decisions.

In pharmaceutical manufacturing, this method to query data structure quickly creates constraints, since questions are often time-bound, documentation-heavy, and tied to inspection readiness—where teams need both an answer and the ability to support it with consistent definitions and traceable records.

With advances in Generative Artificial Intelligence (GenAI) and large language models (LLMs), natural language interaction (NLI) simplifies access and usability by making advanced analytics available through conversation. Databricks Genie represents an emerging way to make analytics feel more conversational while remaining anchored to curated data, definitions, controls and governance.

In practice, inspection-ready answers tend to share a few traits:



Traceable to approved sources



Consistent with agreed definitions and business rules



Reproducible with saved SQL or governed view



Reviewable with clear assumptions, filters and time windows

These traits can help reduce rework and ambiguity while strengthening the narrative that connects KPIs to record-level evidence.





Regulatory compliance in pharmaceutical manufacturing: the data problem

Achieving compliance in pharma manufacturing is complex and distributed across multiple functions and domains, including:



Regulatory compliance



Internal policy compliance



Manufacturing quality assurance and controls



Safety standards



Environmental compliance



Supplier compliance

Large life sciences firms invest significant effort preparing for regulatory inspections by the U.S. Food and Drug Administration (FDA) and other independent organizations. In practice, inspection readiness often hinges on a variety of data types from disparate internal and external data sources. While dashboards can help monitor KPIs and summarize trends, site staff often still spend substantial time manually analyzing underlying data to address specific audit findings, exceptions and deviations ahead of an inspection.

The following tables summarize the major data sources and data types teams typically pull together for a defensible inspection narrative.

MAJOR DATA SOURCES

DESCRIPTION

Batch records	Records each manufacturing run, including ingredients, process parameters, equipment used, deviations, operator identifiers and time logs.
Quality control and laboratory data	Includes analytical test results, in-process testing, product release data, method validation and stability testing records.
Equipment logs	Covers calibration, maintenance, cleaning and performance records for critical equipment and automation systems.
Environmental monitoring	Tracks temperature, humidity, particulate levels, microbiological contamination and other controls in manufacturing and storage areas.
Raw materials source and traceability	Includes supplier certificates, incoming inspection records and batch-level traceability for raw materials.
Personnel training and qualification records	Maintains training completion and competency records for staff involved in manufacturing, quality control and equipment handling.
Audit trails and electronic signatures	Secure, time-stamped logs of changes, approvals and access to electronic records, supporting Title 21 Code of Federal Regulations Part 11 (21 CFR Part 11) requirements.

TYPES OF DATA REQUIRED

DESCRIPTION

Structured data	Numeric batch parameters, test results, log entries, timestamps, sensor readings and system validation results.
Unstructured data	Narrative reports, deviation investigations, incident summaries, corrective action documentation, certificates (e.g., ISO certificates), Good Manufacturing Practice (GMP) documentation, standard operating procedures (SOPs), policy documents, risk assessments and audit reports.
Metadata	Contextual details that enable traceability and compliance verification, such as date/time, operator, equipment and record version.

Leveraging LLMs with retrieval-augmented generation (RAG) can effectively summarize large sets of documents—such as SOPs, policies and prior investigations—when sources are curated and recency is controlled.

However, inspection readiness often requires joining record-level facts across systems (lots, equipment, time windows, approvals) and producing answers that are consistent with governed definitions.

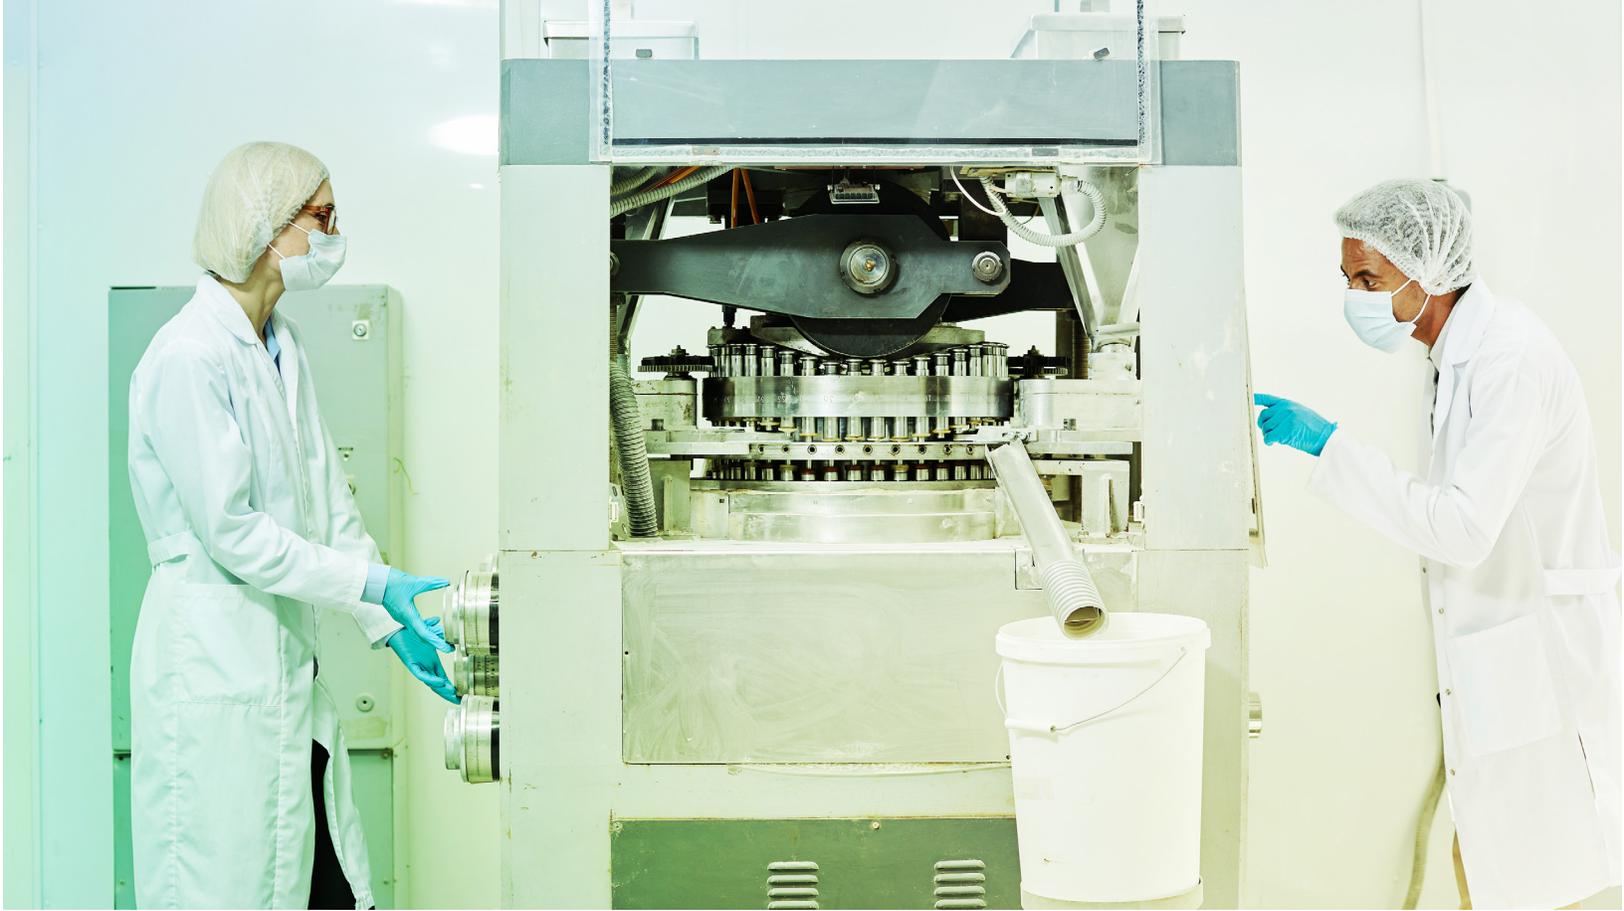
Simply relying on RAG can restrict interactions primarily to unstructured content, such as PDFs and documents, and may not provide a straightforward way to prioritize—or limit—results to the most current information. Organizations need a solution that lets users ask relevant questions and receive responses based on all relevant and available data.

In practice, many teams use both: RAG for narrative-heavy content and Genie for governed, query-backed analytics—linked through shared terminology and controlled sources

A unified interface that enables natural-language exploration of a KPI or data point provides the experience users increasingly expect. With Databricks Genie, teams can explore KPIs in natural language and summarize textual inputs (such as inspector or auditor comments). This reduces time spent assembling and interpreting data and keeps teams focused on inspection preparation.

Additionally, organizations can broaden site staff access to these data assets while maintaining control through centralized governance. That governing experience is anchored in a “Genie space”—a curated, permissioned set of data sources plus metadata and guidance, which give Genie consistent context for answering questions. Teams can then monitor the Genie space for usage patterns and user feedback on accuracy and usefulness, then apply those signals to improve underlying inputs over time.





Implementing Databricks Genie for inspection readiness

Deloitte and Databricks began building an inspection-readiness solution that integrates several interacting components and coordinates calls to multiple models.

Data residing across multiple locations was first consolidated onto a single platform using simple pipelines, then curated into a Genie space with the right data sources, business logic and guardrails. This helps ensure responses remain consistent and inspection-ready.

We can think of a Genie space as an end user's curated, governed analytics environment in Databricks. Here, we can gather relevant data sources, business logic, and instructions—so Genie can deliver accurate, context-aware answers.

HOW DOES A GENIE SPACE WORK?

A Genie space functions by letting users select data sources and optionally add metric views that let them define calculated KPIs commonly used in dashboards, alerts, and other AI functions. The configurations and calculations are created in YAML, registered in Unity Catalog, and can be queried with SQL.

To strengthen the semantic layer, we curated table- and column-level metadata —descriptions, synonyms and representative sample values. We also built value dictionaries that map business jargon to the specific field values Genie should filter on.

Next, common natural-language queries were curated and provided to Genie as usage guidance, paired with the SQL used to derive each answer. This helps Genie align to how end users ask questions—capturing common aliases and domain jargon—and improves responses to similar queries over time.

HOW GENIE CAN LEARN

Genie can also learn from popular queries linked to data assets; however, because many users accessed data primarily through dashboards, the variety of historical queries available for learning was limited.

To maintain trust, we suggest a lightweight review workflow for each Genie space:



A BUSINESS OWNER
(quality/
manufacturing)



A DATA STEWARD



A PLATFORM OWNER
(security/
governance)

Together, they can approve

1. Which sources are in-scope
2. Key definitions and value dictionaries
3. Any promoted “trusted” views or reusable SQL assets

This keeps iteration fast while preserving clarity on what is authoritative.

EXAMPLE OF A VALUE DICTIONARY:

A deviation category field within quality processes of life sciences can be mapped to synonyms such as “exceptions,” “nonconformance,” and “atypical result.” Creating aliases in Genie helps it interpret a variety of questions using any of the synonyms.



Connections matter: Defining relationships across sources

To enable Genie to answer questions that span multiple sources while making the underlying relationships easier for the model and reviewers to interpret, we defined table relationships (for example, one-to-many and one-to-one) and created readable aliases for complex joins.

Through iterative testing, we continuously refined a knowledge store by updating:

- **Instructions to guide response conventions, such as rounding decimals to two places**
- **Example queries with matched SQL for common questions**
- **SQL functions to encapsulate complex or reusable logic**

Finally, we integrated the Genie Conversation API into a custom user interface (UI) that respected the Genie space access controls. The UI included a short FAQ list to help users get started and supported context-preserving follow-ups, allowing users to iterate on questions without restating prior assumptions or filters.

In this approach, Deloitte helps operationalize a repeatable way of working around curation, validation and change control, while Databricks provides the governed platform capabilities, such as Unity Catalog for access control and lineage. This collaboration helps maintain solution trustworthiness while expanding to additional sites, products or investigation patterns.



Inspection-readiness impact

Pharma manufacturers can strengthen quality and compliance by assessing inspection readiness consistently across sites and accelerating investigation workflows. When implemented with curated data and clear guidance, expected outcomes may include:

- **Estimated 15%-18% efficiency gain** in identifying and categorizing risk signals, helping teams expedite preparedness for future audits
- **Estimated 20%-25% reduction in manual errors**, improving accuracy and consistency in inspection preparedness

HOW TO INTERPRET THE ESTIMATES

Efficiency and error-reduction estimates will vary based on baseline data quality, system fragmentation and how consistently sites follow standard investigation playbooks. Organizations typically validate impact through a short pilot that measures cycle time to answer common inspection questions, rework rates and the proportion of answers tied to approved sources and reusable SQL.

Best practices

A conversational layer only helps if it remains fast, consistent and governable as adoption grows. The following practices focus on what most directly drives answer quality and auditability: curation, semantics, performance and access control. These practices can also keep Genie answers accurate, fast and inspection-ready as adoption scales.



Space curation and data preparation

Start each Genie space with a small set of essential tables (five or fewer) and minimal instructions; expand iteratively based on user feedback

Ensure tables and columns in Unity Catalog (Databricks' governed metadata and access-control layer) are clearly named and described; precise annotations improve response accuracy and reduce ambiguity

Keep global context in plain text, and include only a few high-value example SQL queries to “teach” preferred patterns

Optimize performance by keeping examples concise and pruning unused columns; proactively manage token limits to improve response quality and latency



Performance optimization

Encapsulate complex logic in trusted assets (for example, validated views) so Genie can return timely, consistent results

Keep example SQL statements concise and focused to reduce bottlenecks and improve responsiveness



Security and governance

Enforce least-privilege access using Unity Catalog permissions, and use single sign-on (SSO) via an enterprise identity provider so access is role-based and centrally managed

Apply row-/column-level controls, masking where needed for sensitive fields to prevent unintended exposure in responses

Classify and curate approved sources for each space and document what is explicitly out of scope

Regularly audit security posture using Databricks system tables; if needed, stand up a “Security and Trust” Genie space to monitor logs and analyze posture

Establish lightweight change control for adding tables, updating definitions and promoting trusted views

How we measure value

To keep the story grounded, build measurement into the pilot from day one. Many teams baseline the time to answer a small set of priority inspection questions and the amount of rework required, for example follow-up pulls and definition clarifications. They then track whether Genie—implemented on governed Databricks data—reduces that effort while improving repeatability through approved sources and reusable SQL or trusted views.

Turning data into trusted decisions

When Genie spaces are thoughtfully curated and supported with strong data preparation, performance tuning and security controls, teams can move faster without compromising trust. These leading practices help you deliver consistent, auditable answers, reduce rework for data teams and scale self-service analytics with confidence.

A governed conversational experience can help teams spend less time assembling evidence and more time strengthening the narrative that supports quality and compliance. The highest-return starting point is usually one inspection-ready use case, one curated space and a short cycle of testing with site users. From there, organizations can scale confidently by standardizing definitions, promoting trusted logic and operationalizing change control.

AUTHORS



Vikranth Gudala

Managing Director

AI & Data

Deloitte Consulting LLP

vigudala@deloitte.com



Dalveer Rajput

Managing Director

Supply Chain and AI & Data

Deloitte Consulting LLP

drajput@deloitte.com



Deepak Tanna

Sr. Architect

AI & Data

Deloitte Consulting LLP

detanna@deloitte.com

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited (DTTL), its global network of member firms, and their related entities (collectively, the "Deloitte organization"). DTTL (also referred to as "Deloitte Global") and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/about to learn more.

Deloitte provides industry-leading audit and assurance, tax and legal, consulting, financial advisory, and risk advisory services to nearly 90% of the Fortune Global 500® and thousands of private companies. Our people deliver measurable and lasting results that help reinforce public trust in capital markets, enable clients to transform and thrive, and lead the way toward a stronger economy, a more equitable society, and a sustainable world. Building on its 175-plus year history, Deloitte spans more than 150 countries and territories. Learn how Deloitte's approximately 415,000 people worldwide make an impact that matters at www.deloitte.com.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

Copyright © 2026 Deloitte Development LLC. All rights reserved.