Applied Data & AI Engineering in Healthcare

13 times Brillio helped healthcare and life sciences organizations modernize operations, insights, and experiences

brillio

Healthcare organizations are under growing pressure to improve experiences, reduce operational friction, and make better decisions in environments shaped by complexity, regulation, and constant change. Across the enterprise, teams are looking for practical ways to apply AI not as a future ambition, but as a tool for solving real problems today, from onboarding and data management to analytics, translation, compliance, and clinical guidance.

This collection of 13 case studies highlights how Brillio is helping healthcare and life sciences organizations put AI to work in targeted, high-impact ways. Together, these examples show how agentic systems, LLMs, conversational interfaces, and domain-aware automation can improve speed, accuracy, scalability, and user experience across a wide range of business and technology workflows.

Automating provider intake with agent-based orchestration

When a major healthcare organization had to process 800K+ provider records across 10+ formats and 100+ fields, while dealing with duplicate records, missing NPIs/TINs, costly manual mapping and schema checks, and an intake process that affected revenue, the struggle to modernize provider intake became real. The transition toward GraphDB-ready structures added another layer of complexity, requiring changes to mapping, relationships, and validation logic.

To overcome these challenges, the client turned to Brillio. We have designed a unified agent architecture orchestrated through a human-in-the-loop workflow to streamline provider intake at scale.
The solution used three specialized agents: a Format Parser to classify and interpret provider files across varied layouts, a Credentialing Agent to validate credentials against sources such as NPPES and CAQH, and a Validation Agent to detect duplicates, normalize taxonomy data, and resolve inconsistencies before downstream loading.

We also prepared the organization for GraphDB load readiness by automating field inference, credential matching, and validation across structurally inconsistent provider data. Built on Brillio's ADAM framework, it incorporated memory-enabled agent learning and evolving schema resolution to better adapt to changing formats and relationship rules over time. To keep the model practical and scalable, Brillio used a hybrid deployment approach that combined agent-led orchestration with rule-based automation, leveraging the best of existing components instead of forcing a full system replacement.

Business Impact

80% reduction
in manual parsing

60% reduction
in handoffs

5x
operations productivity gain

Improved
CMS audit readiness

95%+
reduction in stale records

Scaled
across 67,000+ physicians and 263+ hospitals

Reinventing provider data management with a team of agents

A leading health insurer was facing growing provider data management challenges as it worked to modernize a legacy PDM environment. The organization was dealing with 250K+ provider records across 10+ formats and 100+ fields, along with duplicate entries, missing NPIs/TINs, legacy data quality issues, and high manual effort tied to mapping, schema checks, and rework. These issues were slowing onboarding, increasing CMS compliance risk, and making migration to a modern provider data model harder to scale.

We've addressed this by building a unified agent architecture spanning the lifecycle of both PLM and PDM, using a team of 13+ agents to support a broader, multi-purpose transformation model that could extend to other domains over time. Rather than treating provider data modernization as a one-time migration exercise, the solution was designed to work across evolving intake, mapping, validation, and orchestration needs in a more flexible way.

The solution design focused on three priorities: driving user adoption, enabling a composable AI architecture, and maximizing ROI. To support this, Brillio leveraged its ADAM framework with memory-enabled agent learning and evolving schema resolution, allowing the system to adapt as data structures changed. The deployment also followed a cost-conscious hybrid model, combining agent-led workflows with rule-based automation to make better use of existing components while improving scalability, resilience, and transformation readiness.

Business Impact

80% reduction
in manual parsing

80% automated
ingestion of rosters

60% reduction
in handoffs

5x
operations productivity gain

30% reduction
in engineering effort

Improved
CMS audit readiness

95%+
reduction in stale records

Scaled
across 100,000+ practitioners, 8,000+ provider groups, and 2 million members

Accelerating financial decision-making with an LLM-powered conversational agent

With financial analysts and developers having to write SQL queries by hand to extract data for reports, which increased effort, introduced inconsistency, and delayed decision-making, this major life sciences company turned to Brillio to help them overcome the slow, manual financial insight generation. As reporting complexity grew, manual query creation became a bottleneck that limited the speed and accessibility of business insight.

We have developed an LLM-powered conversational agent integrated with Cortex Analyst on Snowflake. The solution enabled users to ask business questions in natural language, automatically generate structured SQL queries based on rules and schema logic, and receive insights on trends, KPIs, and anomalies without depending on manual coding. It also supported LLM-powered chart suggestions and query reasoning so users could better understand how results were derived.

By reducing the effort required to translate business questions into analytical outputs, the solution helped move finance teams closer to faster, more intuitive decision support. It also showed how conversational analytics can improve accessibility without compromising analytical rigor.

Business Impact

80% reduction
in manual effort to create SQL queries

75% reduction
in errors during insight generation

66% faster
decision cycles

Automating contract data retrieval with Cortex LLM

A major healthcare organization was struggling to retrieve and reconcile contract intelligence spread across multiple systems. Critical information tied to speaker and consulting programs, including contracts, payment terms, event details, compensation, and related performance data, was fragmented across structured and unstructured sources, making retrieval slow, analysis cumbersome, and insight generation difficult for legal, marketing, compliance, and leadership stakeholders.

We have developed chatbots leveraging Cortex Search and Cortex LLM to automate extraction of contract metadata, including key clauses such as payment terms, deliverables, timelines, compensation, terms of agreement, duration, and speaker metadata. The solution also used Cortex Analyst to retrieve and analyze related tabular data, such as payment history, program milestones, and speaker performance, allowing users to more quickly retrieve, compare, and analyze contract terms across multiple sources.

The solution searched, analyzed, and indexed unstructured data such as speaker agreements and SOWs stored in Snowflake, while also constructing optimized SQL queries based on business rules and Snowflake table schemas to retrieve relevant structured data. By analyzing both structured and unstructured sources together, it enabled users to generate more relevant insights on program success rates, speaker engagement, and ROI of consulting contracts.

Business Impact

80% reduction
in retrieval effort for contract terms through automation

70–80% faster
insight delivery

75% cost savings
through automated data gathering and reconciliation

Moving from static reports to dynamic business guidance with conversational agents

Limited by static reporting experiences that made it difficult for business users to move beyond fixed dashboard views, this leading life sciences enterprise partnered with Brillio for help. Their existing Qlik dashboards constrained users to static visualizations, offered limited cross-report navigation, and did not support real-time exploration of unstructured data. As a result, generating dynamic insights still required more manual effort than the business could afford.

We have integrated an intelligent Snowflake-based conversational GenAI chatbot within the existing Qlik dashboard environment to streamline insight generation for field representatives and other business users. The solution was designed to enhance user interaction, simplify data access, boost field productivity, and support faster insight-driven decision-making through natural language queries. At the core of the experience, Cortex Agent acted as an intelligent router to analyze incoming queries and direct them to the right service based on whether the request involved structured or unstructured data.

The solution used Cortex Analyst to interpret natural language queries against structured POA data in Snowflake, while Cortex Search processed unstructured sources such as FAQs, historical reports, and guidelines. It also enabled multi-turn responses for ambiguous questions, allowed users to navigate across reports and dashboard sections conversationally, and leveraged the structured data sources already powering Qlik visualizations to deliver direct answers. To strengthen the experience over time, the solution incorporated user feedback and usage metrics to improve response quality and expand capabilities, while maintaining a consistent omnichannel experience across laptops, tablets, and mobile devices.

Business Impact

75–85% reduction
in time to insight generation

40% increase
in actionable insights for decision-making

56%
cost savings

Extracting actionable insight from coaching data with a Cortex LLM chatbot

A leading life sciences company was struggling to extract value from coaching and feedback data scattered across both structured and unstructured formats. Post-session reviews, meeting notes, competency metadata, and related inputs contained valuable signals, but manual analysis made it difficult to surface timely patterns, sentiment, and actionable recommendations. That slowed reporting and limited how effectively the organization could use coaching data in decision-making.

Through an LLM-powered chatbot integrated with Cortex Search and Snowflake, we have automated insight extraction from coaching data for the client. The solution enabled users to query both structured and unstructured data in natural language, construct optimized SQL for structured sources, and combine that with indexed feedback content to surface trends, performance patterns, and feedback themes. This created a more unified view across notes, reviews, competency metadata, and related inputs.

The result was a more efficient and scalable way to turn coaching data into usable business insight. It also showed how conversational access can improve the value of data that would otherwise remain buried in text-heavy formats.

Business Impact

80% accuracy
in sentiment analysis

60 - 80%
of extracted insights used in decision-making

30% cost savings
through faster time to insight

From 4 weeks to 1 hour: Automating static data preparation with OpenAI

A major healthcare company was spending more than a month on static data preparation during enhancement cycles, turning a routine activity into a major delivery bottleneck. The work relied heavily on manual methods, which increased the risk of human error, consumed significant developer effort, and slowed both configuration updates and downstream release timelines. In practice, a process that should have taken hours was stretching to 4 weeks, with broader development cycles extending to 8 weeks.

We've addressed this by using OpenAI and Azure OpenAI to automate static data preparation workflows. The solution generated JSON-based static data files that could be prepopulated and surfaced through a web wizard, allowing users to review, edit, delete, or extend them as needed. AI-generated database scripts could then be uploaded for final review and approval, reducing the burden on development teams while maintaining control.

This significantly compressed a process that previously took weeks and made static data preparation far easier to operationalize at scale. It also demonstrated how AI can remove friction from a highly repetitive SDLC activity with immediate measurable returns.

Business Impact

50% reduction
in manual effort

99.85% reduction
in static data preparation time, from 4 weeks to 1 hour

50% reduction
in development cycle time, from 8 weeks to 4 weeks

Optimizing security and vulnerability remediation with LLMs

Trying to accelerate cloud security remediation after the scanning environment had already identified thousands of vulnerabilities and converted them into VITs in ServiceNow with defined SLAs was a challenge for this leading biopharmaceutical company. While discovery was working, remediation still required teams to navigate multiple platforms and disparate information sources to inspect code, patch packages and runtime environments, and close issues quickly enough. The result was a manual, delay-prone process that created operational drag in a high-stakes security workflow.

We have addressed this for our client by integrating an LLM workflow into the remediation process. Using Wiz reports, Lambda source code, AWS Bedrock, and supporting AWS services, the solution analyzed issues, generated outputs, and helped automate source code updates, versioning, and deployment steps. The workflow also included safeguards such as skip logic and chunking to manage token limits and reduce failures.

This created a faster and more scalable remediation model while reinforcing the need for explainability, governance, and confidence-based controls in LLM-enabled security operations.

Business Impact

80.9%
overall success rate

80% reduction
in manual effort through LLM enablement

5x
operations productivity gain

560
executions completed

356
VITs targeted and 90 VITs closed

453
total updates applied, including 141 package-level and 312 runtime-environment updates

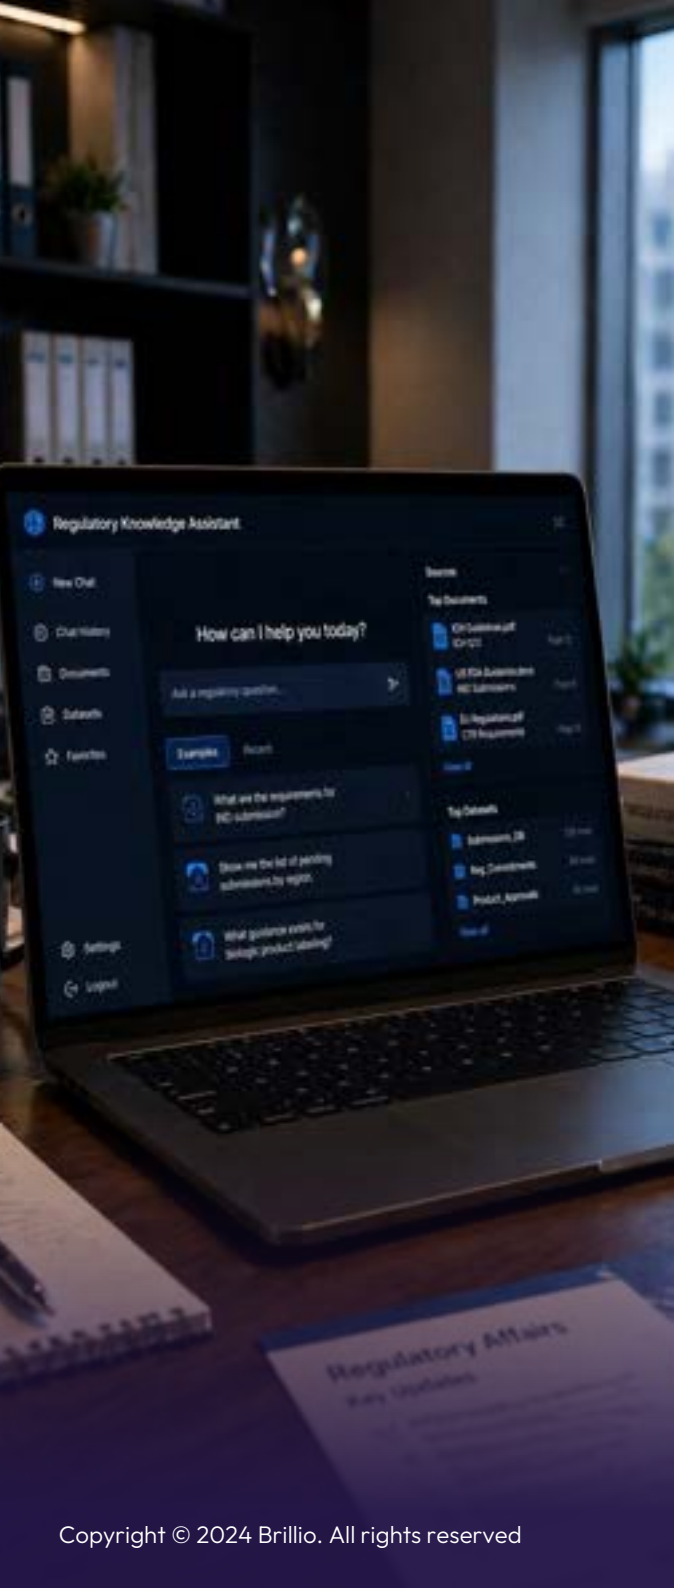

Centralizing knowledge management with a conversational Q&A AI agent

A major biopharmaceutical company's regulatory business team needed a faster way to answer specific business questions using enterprise data stored across both documents and tabular datasets. Instead of having a single interactive interface, users often had to rely on slower request-driven processes and navigate fragmented information sources to find what they needed. That made knowledge access inefficient and limited the speed at which the business could retrieve critical answers.

We built a conversational Q&A agent designed to answer business questions across both document-based and tabular enterprise data. For document workflows, the solution used AWS OpenSearch as the retriever to surface relevant top-k content, then passed that output to Azure OpenAI to generate natural-language responses with prompt engineering. For tabular workflows, the solution used Azure OpenAI to generate SQL queries based on the defined schema, sent those queries to a SQL engine to retrieve the required output, and then used an additional Azure OpenAI layer to convert the result into a natural-language response. LangChain was used to add the conversational layer across both flows.

To connect these experiences, Brillio developed a multi-classification model that automatically identified whether a user query was structured or unstructured and routed it to the appropriate workflow. This gave the regulatory business team a single interactive interface for retrieving answers from mixed data sources, replacing a slower request-driven process with more direct, self-service access to enterprise knowledge.

Business Impact

Created a one-stop tool
to access information from documents and tabular datasets

Enabled users to access
required information instantly instead of following a lengthy request process

Achieving up to 60% cost reductions through LLM-enabled content translation

Struggling to translate patient-facing portals, EMRs, medical forms, and instructional content with the speed, precision, and compliance required in regulated healthcare environments, this major biopharma organization turned to Brillio for help. Manual translation was slow, error-prone, and costly, especially for medical jargon, while evolving health content such as advisories and dosage information made real-time localization increasingly important. At the same time, regulatory requirements such as HIPAA and MDR demanded precise translation, auditability, and secure handling of PHI across fragmented systems.

We tackled this challenge by building an LLM-enabled medical translation approach that combined fine-tuned medical LLMs with domain-focused NMT models to improve biomedical translation quality while preserving terminology accuracy. The solution used clinical ontologies such as FHIR, SNOMED CT, and ICD-10 to inject medical context and improve consistency, while lightweight expert small-language models supported critical terms through an “LLMs-in-the-loop” approach. A human-in-the-loop QA workflow with PHI auditing was integrated using HIPAA-compliant review controls, and the translation API was designed for secure deployment across EHRs, telehealth platforms, CMS environments, and patient-facing portals.

This created a more scalable and compliant localization model that improved consistency, accelerated translation turnaround, and reduced the operational burden of medical content translation. It also reinforced that effective healthcare translation depends not just on language generation, but on domain tuning, ontology grounding, secure workflow design, and structured human review.

Business Impact

30% improvement
in patient engagement

50% faster
time-to-market for global product launches

40–60% reduction
in localization costs

90%+
compliance accuracy

Optimizing SQL performance, analysis, and integration through AI-led engineering

A major HLS enterprise was facing performance and optimization challenges during an Oracle-to-PostgreSQL migration, where differences in architecture, query behavior, and tuning patterns created new bottlenecks. Query improvements required extensive database team involvement, repeated review cycles, and deep understanding of both the data model and business context. As PostgreSQL-specific optimization needs became more apparent, the organization faced growing dependency on scarce DB expertise to analyze execution plans, tune queries, and close knowledge gaps introduced by the migration.

We've addressed this through an AI-led engineering approach built to improve query analysis, optimization, and integration in PostgreSQL environments. The solution used Claude AI to analyze execution plans and identify PostgreSQL-specific bottlenecks, while Cursor AI was connected directly to the database through MCP PostgreSQL server integration to provide real-time context without depending on manual data model documentation. It also supported automated optimization generation tailored to PostgreSQL architecture, plain-English query generation for accessibility, and EXPLAIN plan analysis to recommend practical improvements in index usage, join ordering, filter reordering, and performance-oriented query restructuring.

This helped accelerate the query optimization workflow, improve engineering productivity, and reduce dependence on the database team for routine tuning and SQL generation tasks. It also showed that AI can add real engineering value when paired with database-team validation, strong business context, and feedback loops between suggested optimizations and actual performance outcomes.

Business Impact

35% improvement
in team productivity

Accelerated
query optimization workflow

Reduced
DB team dependency

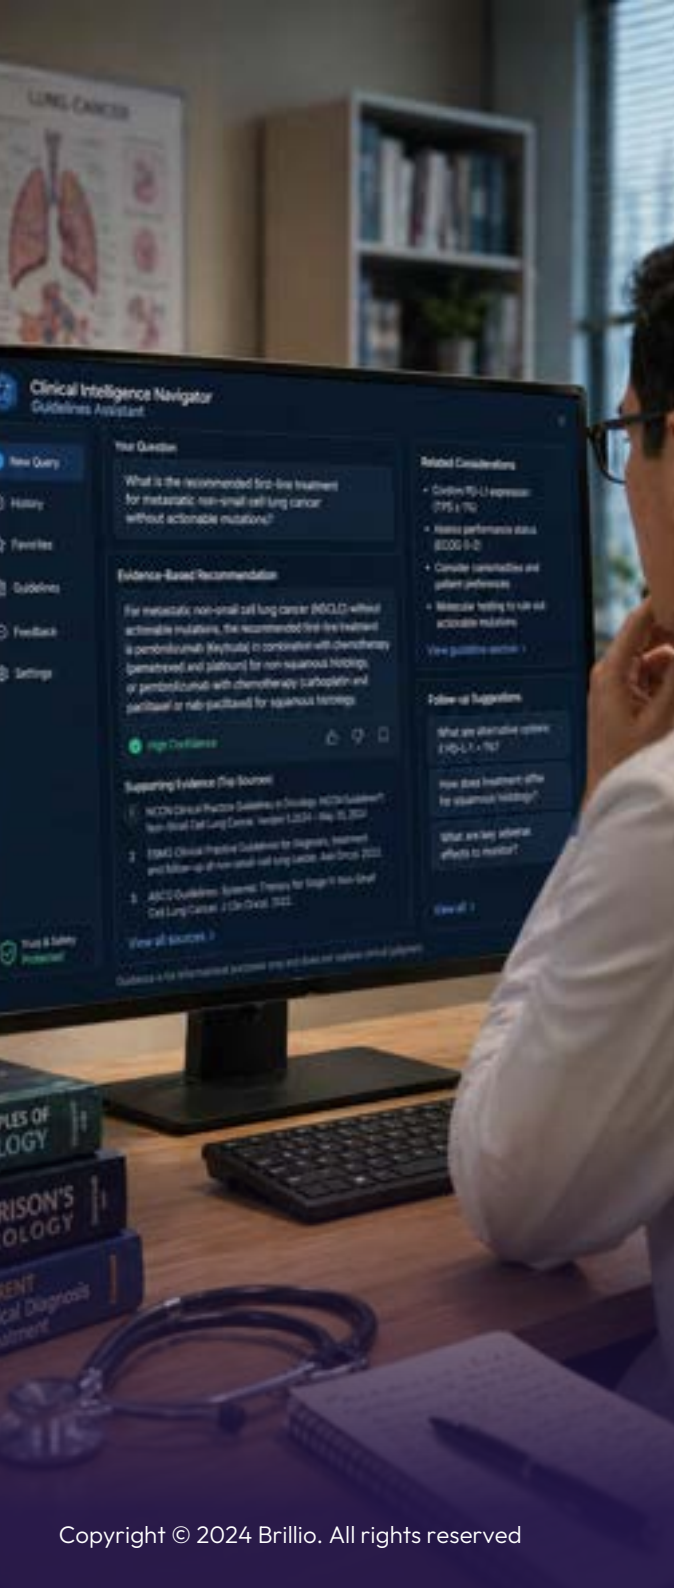

Delivering evidence-based treatment recommendations through a clinical intelligence navigator

Healthcare providers often struggle to navigate complex and evolving clinical guidelines, leading to variability in care and delays in decision-making. When oncology guidance is difficult to discover, interpret, or apply at the point of need, clinicians may have less confidence that they are accessing the most relevant recommendations quickly enough to support consistent care delivery.

To solve this challenge, we've built an AI-powered Guidelines Assistant on GCP to improve the discoverability and accessibility of clinical oncology guidelines for HCPs. The solution used a fused RAG workflow and APIs to handle user queries, retrieve relevant documents, and generate grounded answers based on underlying guidance. To improve safety and trust, it incorporated safety and relevance checks to identify harmful categories and reduce the risk of PHI/PII exposure in member queries. The experience was further strengthened through synonym matching for improved retrieval, starter and follow-up workflows tailored to clinical guidelines, and a robust continuous evaluation architecture with feedback loops and IaC to support efficient improvement over time.

The result was a more reliable and confidence-inspiring experience for clinical guideline delivery, helping HCPs access more relevant recommendations with greater trust in answer quality. By combining grounded retrieval, safety controls, and continuous evaluation, the solution improved both factual performance and the usability of evidence-based treatment guidance.

Business Impact

98.75% factuality and completeness
improving confidence in clinical guideline delivery

Quality scores improved
from **78.79% to 98.75%** across iterations

ABOUT BRILLIO

Brillio is one of the fastest growing digital technology service providers and the partner of choice for many Fortune 1000 companies seeking to turn disruption into a competitive advantage through innovative digital adoption. We help clients harness the transformative potential of the four superpowers of technology: cloud computing, Internet of Things (IoT), artificial intelligence (AI) and mobility. Born digital in 2014, we apply our expertise in customer experience solutions, data analytics and AI, digital infrastructure and security, and platform and product engineering to help clients quickly innovate for growth, create digital products, build service platforms, and drive smarter, data-driven performance. With 17 locations across the U.S., the UK, Romania, Canada, Mexico, and India, our growing global workforce of nearly 6,000 Brillians blends the latest technology and design thinking with digital fluency to solve complex business problems and drive competitive differentiation for our clients. Brillio has been certified by Great Place to Work since 2021.

https://www.brillio.com/
Contact Us: info@brillio.com

brillio