VIEW POINT

SECURE KNOWLEDGE RETRIEVAL IN LIFE SCIENCES: UNLOCKING ENTERPRISE KNOWLEDGE THROUGH AI WHILE PRESERVING COMPLIANCE AND GOVERNANCE

Infosys | CONSULTING

Executive Summary

Life Sciences organizations hold vast regulated knowledge across Regulatory Affairs, Quality, Clinical, Medical Affairs, Pharmacovigilance, Manufacturing, and R&D. Yet this information is often fragmented across platforms, records, metadata, and documents, making it difficult for users to quickly find trusted answers for operational and compliance decisions.

Generative AI can transform how users access enterprise knowledge, but regulated environments require more than conversational search. Solutions must preserve security, role-based access, traceability, auditability, and confidence in the accuracy of responses.

Secure Retrieval-Augmented Generation provides a practical path forward by grounding AI responses in authorized enterprise information. This paper explores how Life Sciences organizations can use secure, governed AI-powered retrieval to unlock knowledge, reduce manual effort, and scale trusted AI adoption across functions.

Infosys' implementation experience across Regulatory Affairs and Quality Management demonstrates that secure AI-powered retrieval can reduce search effort, improve access to evidence, and support faster, source-backed decision-making in regulated Life Sciences environments.

The Knowledge Discovery Challenge in Life Sciences

Life Sciences functions operate in highly specialized and regulated environments, with Regulatory, Quality, Clinical, Medical Affairs, Pharmacovigilance, Manufacturing, and R&D teams managing large volumes of documents, records, metadata, workflows, correspondence, and compliance evidence.

This challenge is becoming more acute as AI adoption expands across Quality Assurance and Regulatory Affairs, where industry commentary highlights the opportunity to improve efficiency, consistency, insight generation, and compliance readiness across submissions, post-market activities, and regulated operations. [1]

Regulatory and Quality users, for example, often need quick answers on submission commitments, version changes, supporting evidence, applicable SOPs, prior CAPAs, and related investigation history.

In most cases, this information already exists within enterprise platforms; the challenge is locating, interpreting, and connecting it efficiently enough to answer business questions.

The issue is no longer information availability; it is information accessibility.

Why Traditional Approaches Fall Short

Organizations have traditionally relied on search engines, reports, dashboards, metadata filters, and subject matter experts to access enterprise knowledge.

While effective for locating records, these approaches often struggle to provide contextual answers.

Approach	Limitation
Keyword Search	Returns records rather than business answers
Metadata Search	Requires understanding of underlying data structures
Reports and Dashboards	Limited flexibility for exploratory questions
SME Dependency	Creates bottlenecks and knowledge silos
Knowledge Repositories	Difficult to maintain and scale

As information volumes continue to grow, users spend increasing amounts of time searching, validating, and synthesizing information before acting. This creates a productivity challenge across multiple Life Sciences functions.

At the same time, generative AI is increasingly viewed as a major productivity lever for the pharmaceutical and medical products industry, with external estimates suggesting potential annual economic value of $60 billion to $110 billion across research, development, operations, commercial, and medical affairs domains. [2]

The Life Sciences AI Opportunity

Generative AI has fundamentally changed expectations around information access. Users increasingly expect to interact with enterprise systems using natural language and receive direct, contextual answers to business questions.

Examples include:

- What commitments are currently open for Product X?
- What were the key changes introduced in the latest submission?
- Which approved procedures are relevant to this quality event?
- What historical decisions were made in similar situations?

However, AI solutions in Life Sciences face unique constraints. Any AI solution must ensure:

- Security of regulated information
- Role-based access control
- Traceability of responses
- Auditability of user interactions
- Alignment with compliance requirements
- Confidence in response accuracy

Generic AI solutions often struggle to meet these expectations because they operate outside the governance framework of regulated business applications. As a result, trust remains one of the primary barriers to enterprise-scale AI adoption.

Regulatory agencies are also recognizing AI's potential to improve regulatory assessment, manufacturing, post-market surveillance, and knowledge management, while emphasizing the need to understand AI's limitations, lifecycle implications, and responsible use in FDA-regulated product contexts. [3]

Secure RAG: Enabling Trusted AI-Powered Knowledge Retrieval

Retrieval-Augmented Generation (RAG) offers a practical mechanism for balancing usability and governance.

Rather than solely relying on the information embedded within a language model, RAG retrieves relevant enterprise information at query time and uses it as grounding context for response generation.

A typical workflow involves:

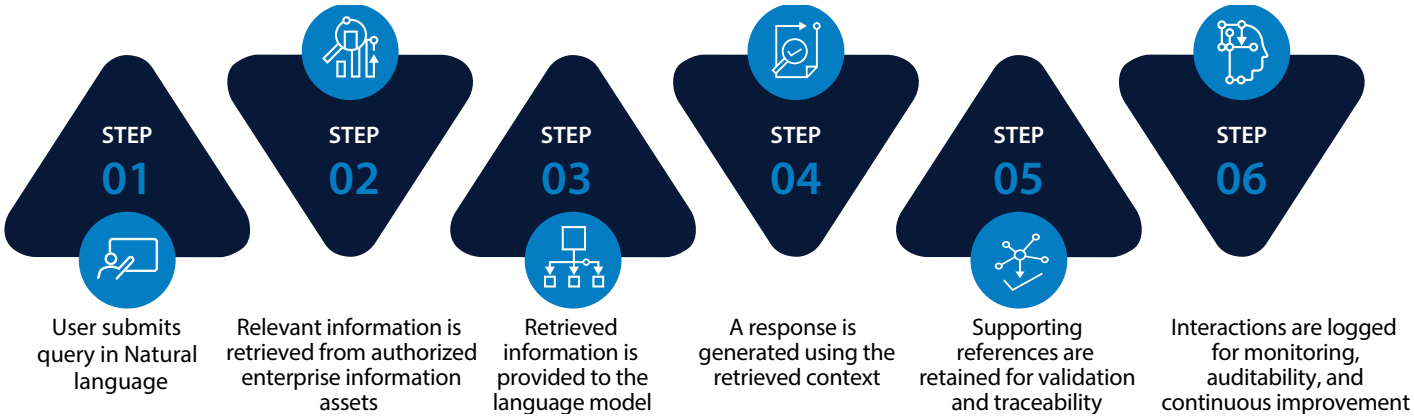

This approach delivers several benefits:

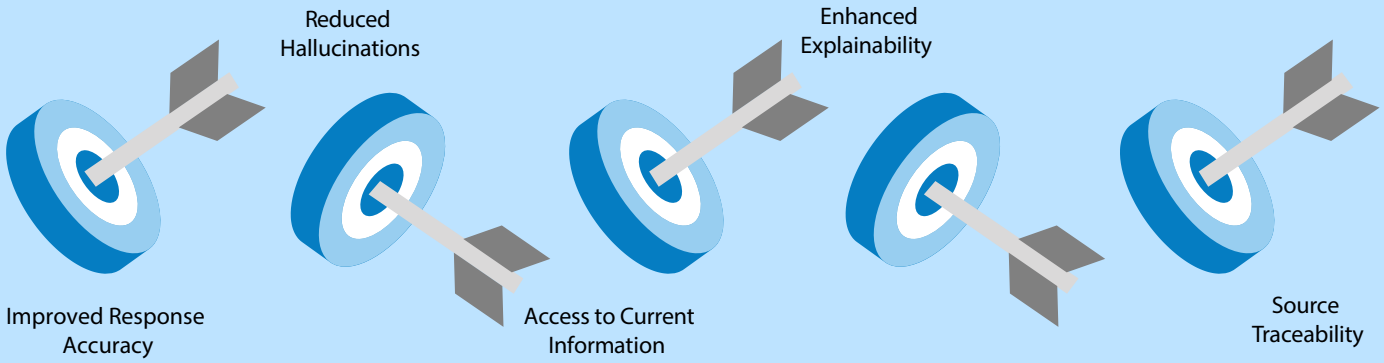

More importantly, it enables organizations to leverage Generative AI while maintaining governance controls.

Accelerating Adoption Through Reusable AI Foundations

While many organizations are exploring AI-powered knowledge retrieval, deployments are often built as standalone solutions tied to specific platforms, data sources, or user interfaces. This can increase implementation effort and limit scalability.

A more sustainable approach is to establish a reusable AI foundation that provides knowledge retrieval capabilities as a service. Such a foundation can be configured to work with different governed platforms and integrated with existing user interfaces, enabling organizations to rapidly deploy function-specific AI assistants without rebuilding core capabilities for every use case.

With this model, organizations can ***"deploy once and configure many."*** The same foundation can be adapted across Regulatory, Quality, Clinical, Medical Affairs, Pharmacovigilance, Manufacturing, and R&D use cases, while preserving consistent security and governance controls.

For business leaders, this means faster access to trusted insights and improved productivity. For IT leaders, it means a scalable architecture that avoids fragmented pilots and supports enterprise adoption. For compliance and quality stakeholders, it means AI can be introduced without compromising traceability, access control, or audit readiness.

This reusable model is especially important in Life Sciences because the long-term value of AI will not come from one isolated assistant. It will come from a governed ecosystem of AI capabilities that can safely unlock enterprise knowledge across functions.

From Concept to Practice: Infosys Experience

Infosys has been actively working with Life Sciences organizations to operationalize AI-powered knowledge retrieval within regulated environments.

In Regulatory Affairs, the solution enabled natural language access to regulatory content and structured metadata. Users could identify submission information, retrieve supporting evidence, understand document changes, and prepare health authority responses faster. Source-backed answers helped reduce manual effort and improve confidence.

In Quality Management, Infosys enabled conversational retrieval across deviation data and related quality records. Users could query deviation trends, linked CAPAs, implemented actions, and recurrence of similar issues. The solution preserved traceability to underlying QMS records while simplifying access to insights.

Beyond individual implementations, Infosys has developed reusable AI foundations that enable rapid development of knowledge retrieval capabilities across different business functions and platforms. This approach reduces implementation effort, accelerates time-to-value, and provides a consistent framework for governance, traceability and security.

The capabilities demonstrated can be directly mapped to the measurable benefits delivered:

Capability demonstrated	Measurable benefit
Grounded responses with traceability and source references	100% citation-backed responses
Context-aware question answering over regulated information assets	60–70% reduction in information retrieval time
Retrieval from structured and unstructured enterprise information	Over 95% accuracy for structured data queries
Natural language retrieval for Regulatory Affairs use cases	2–4 hours saved in health authority query response preparation
Integration with existing enterprise platforms and user interfaces	Reduced dependency on manual search and SME support
Reusable deployment patterns for new use cases	Faster reuse of governed enterprise knowledge
Role-based access controls, user entitlements, governance, and auditability	Ensures benefits are delivered within trusted compliance and security boundaries

These implementations demonstrate that secure AI-powered retrieval is no longer a future state aspiration. It is an achievable capability that can be deployed today within regulated Life Sciences environments.

Conclusion

Life Sciences organizations do not lack information. They lack fast, trusted, and governed access to the knowledge already embedded in their enterprise platforms.

Secure RAG enables organizations to transform regulated content, records, and metadata into contextual, citation-backed answers while preserving the controls required for compliance, security, and auditability. When implemented with the right architecture and governance, AI-powered retrieval can reduce manual effort, improve decision speed, and increase the value of existing platforms such as Veeva Regulatory and Veeva QMS.

Infosys' implementation experience shows that this capability is achievable today. With reusable AI foundations, Life Sciences organizations can move beyond isolated pilots and build scalable, governed knowledge retrieval capabilities across functions.

The future of AI in Life Sciences will not be defined by generic assistants alone. It will be defined by trusted, compliant, and domain-aware AI solutions that help users convert regulated enterprise information into actionable knowledge.

References

1. https://www.iqvia.com/-/media/iqvia/pdfs/library/white-papers/2025/ai-in-life-science-quality-management-and-regulatory-affairs.pdf
2. https://www.mckinsey.com/industries/life-sciences/our-insights/generative-ai-in-the-pharmaceutical-industry-moving-from-hype-to-reality
3. https://www.fda.gov/science-research/focus-areas-regulatory-science-report/focus-area-artificial-intelligence

Authors

Saarthak Gupta

Senior Consultant, Infosys Consulting

Saarthak is a Senior Consultant in Infosys Consulting's Life Sciences practice within the LS Data Transformation Team with more than 7 years of professional experience. He has experience in Life Sciences, Healthcare, Information Technology and Ed-Tech industries. He has worked in multiple engagements on Data Analytics, Data Migration and Transformation, Data Quality Management, Reference Data Management, Net Revenue Management and AI Implementation.

Hemant Kumar

Consultant, Infosys Consulting

Consultant in Infosys Consulting's Life Sciences practice with professional experience in Information Technology and Supply chain. Has 4+ years of experience in Application Development, Agile and Scrum methodologies, Client Communication, Supply chain management, SQL, Power Platform tools, Data modelling, Python, Advanced excel etc.

Naveen NirmalKumar Devaraj

Senior Principal Consultant, Infosys Consulting

A Principal consultant with depth of domain & IT experience specializing in R&D Transformation and Capability Building. I partner with pharmaceutical and life sciences organizations to drive innovation, operational excellence, and sustainable growth. With a deep understanding of the complexities in drug development, clinical research, pharmacovigilance, and regulatory affairs, I help organizations modernize their R&D processes and build robust capabilities that align with evolving industry standards and technological advancements.

ABOUT INFOSYS CONSULTING

Infosys Consulting is a next-generation consulting partner that bridges strategy and execution. With an AI-first mindset, deep industry knowledge, and the combined strengths of business and technology consulting, it helps enterprises turn bold vision into tangible outcomes, faster, smarter, and at scale. Infosys Consulting is helping some of the world's most recognizable brands transform and innovate. Our consultants are industry experts that lead complex change agendas driven by disruptive technology. With offices in 20 countries and backed by the power of the global Infosys brand, our teams help the C-suite navigate today's digital landscape to win market share and create shareholder value for lasting competitive advantage.

For more information, contact consulting@infosys.com

Infosys® | CONSULTING