Analyst Brief

Generative Ai Usage in Regulated Industries as a Strategic Lever for Fractional Technology Leaders

Christopher Richardson
Principal Analyst, Galson Research

Background

Generative Ai continues to present compelling opportunities for integration into day-to-day workflows, yet many organizations still struggle to move past concerns related to cost, risk, and clarity of use case. This analyst brief focuses on a specific, high value application of Generative Ai in regulated industries and translates that experience into practical lessons you can use as a strategic lever in your own work.

Regulated industries such as healthcare, finance, energy and utilities, transportation, food and drugs, telecommunications, gambling, manufacturing, and aerospace are distinct because their regulators are responsible for protecting the public interest while also defining the standards that products and services must meet to remain in compliance. This dual role creates an environment where documentation, traceability, and repeatability are not simply operational preferences, but formal expectations.

As part of ongoing compliance efforts, organizations in these sectors must complete official forms, submit narratives, and maintain evidence at multiple stages of product development and operations. In many cases, compliance also requires well documented processes or systems that can be understood and evaluated by external reviewers. The resulting documentation is often repetitive, highly structured, and formulaic in its responses, which makes it an ideal candidate for Generative Ai support that can maintain consistency while reducing manual drafting effort.

The case study that follows illustrates how this pattern can translate into real outcomes. By targeting a narrow, repeatable documentation domain, the organization was able to save money, reduce risk, and move products through regulatory milestones more quickly, without eroding the level of oversight expected in a regulated environment.

Case Study

Industry:

Medical Device Manufacturing

Respondent:

Fractional COO/CTO, focused on startups

Problem:

Overwhelm with regulatory compliance documentation.

The "Old" Way

The medical device manufacturing industry operates at the intersection of healthcare and manufacturing, both of which are regulated environments with significant compliance obligations. For smaller companies, the traditional approach to compliance often depends on outsourcing substantial portions of the work to specialized consultants because the volume of required documentation is so high.

In this model, internal teams gather and present data to consultants, who then complete the formal compliance activities, while a member of the leadership team reviews the outputs for quality before multiple rounds of rRevision and eventual submission to regulatory agencies. Much of the underlying documentation relies on standard responses with only minor adjustments from one submission to the next, which underscores how repetitive these tasks can be.

Although this consultant-driven model can be effective, it does not scale well as the organization grows or as product complexity increases. It introduces significant time and budget pressures, since consultants are costly and each new document origination effort can require days or even weeks to move from initial draft to a version that is ready to send.

Trial and Error

Generative Ai, delivered through a chatbot interface, was introduced into the workflow to generate required narratives and populate forms with relevant information, with the intent of reducing or even eliminating reliance on external consultants. While the initial concept appeared straightforward and attractive, the limitations of this approach became apparent very quickly.

01

The first challenge involved constructing prompts that contained enough context and specificity to produce drafts that were truly usable. Teams discovered that without careful prompt design, outputs were either too generic or incomplete, which shifted effort from drafting to repeated prompt refinement.

02

The second challenge related to the consistency and reliability of the results. Hallucination rates were unacceptably high for the volume and criticality of the documents being produced, which made it difficult to trust the chatbot as a dependable part of the compliance workflow.

03

The third challenge centered on uncertainty about the underlying data flows that supported the Generative Ai tool. There was uncertainty about how prompted data was stored, whether it needed additional protection, and what it would mean for potentially sensitive outputs to be incorporated into systems that might later become part of public facing databases.

Solution

Since the chatbot-based solution showed promise, the team chose not to explore other access patterns for Generative Ai at that stage. Instead, they upgraded to a team license that provided SOC 2 Type 2 compliance as well as encryption of data at rest using AES-256 and in transit using TLS 1.2, which helped address initial data security concerns. Additional configuration within the chatbot environment was still required to ensure that customer and regulatory data would not be used to train future models, which was an important consideration for a regulated context.

The next hurdle involved improving output quality by strengthening the context available to the model. The team began to design longer and more detailed custom prompts that could be stored, reused, and refined, and they complemented this work by introducing a RAG pipeline that brought relevant reference material directly into the generation process.

RAG, or Retrieval Augmented Generation, is a technique that allows a language model to consult information outside of its original training set in order to deliver more specific and relevant answers to user queries.

This combined approach required more upfront effort and carried a higher operating cost than a simple chatbot configuration, but the return on investment was clear. The workflow became faster and more accurate while also reducing costs relative to the previous consultant-heavy model, in large part because a human remained in the loop to perform quality control at every stage. If a hallucination appeared in an early step, that output was corrected or rejected before it could influence downstream work, which limited risk while preserving the speed benefits of Generative Ai.

A persistent challenge in professional use of Generative Ai is the ability to verify outputs against reliable sources through an automated or semi-automated process. In regulated industries, humans already serve as the first line of quality control, which makes it natural to introduce Generative Ai as a tool that accelerates document creation while still relying on expert review for validation. This context raises an important strategic question: does a Generative Ai process truly need to be fully automated to deliver effective and efficient results at scale?

Given the cost of tokens and the potential complexity of highly automated or Agentic Ai workflows, a more pragmatic strategy is to use Generative Ai as an enhancement to existing workflows, rather than as a complete replacement. In this model, Generative Ai serves as a force multiplier for subject matter experts, who retain accountability for final outputs while benefiting from improved speed, consistency, and access to relevant information.

Recommendation

Our case study demonstrates that Generative Ai delivers the strongest value in regulated environments when it is treated as an enhancement to human-led workflows rather than a fully autonomous replacement. Fractional technology leaders should prioritize architectures and operating models that preserve human quality control while using Generative Ai to compress drafting time, standardize language, and improve consistency across required regulatory artifacts.

First, localize and harden the large language model footprint for the most sensitive and repetitive work. For document types where structure and content change slowly over time, an offline, open source, locally hosted model provides material security advantages by keeping proprietary data and trade secrets inside the organization's own perimeter. This closed-loop design not only reduces the risk of unintentional data leakage into third party systems, it also creates a stable environment where prompts, templates, and responses can be iterated and versioned with greater predictability.

Second, pair the localized model with a well-governed RAG pipeline built on a vector database and a fit for purpose embedding model. Curating a vetted corpus of SOPs, historical submissions, regulatory guidance, and prior approvals gives the model access to the same reference surface that human experts rely on, but with far greater retrieval speed. When implemented correctly, this allows the Generative Ai system to generate drafts that are both context aware and aligned with current interpretations of the regulatory landscape, while making it easier for reviewers to trace statements back to authoritative source documents.

Third, formalize a human in the loop operating pattern that mirrors existing compliance review stages rather than trying to bypass them. Each step in the document lifecycle, including data ingestion, initial drafting, revision, and final approval, should have explicit Generative Ai assist roles and corresponding human sign off, with clear criteria for when an output can proceed or must be reworked. This preserves the accountability regulators expect, while allowing teams to benefit from speed and scale advantages where they are most defensible.

Finally, Fractional Technology Leaders should treat this stack of localized LLM, governed RAG, and structured human oversight as a reusable pattern across multiple clients and regulated domains. By standardizing the architecture and adapting only the domain corpus, leaders can reduce implementation time, simplify risk explanations for boards and regulators, and position themselves as strategic partners in modernizing compliance heavy workflows.

Conclusion

The experience of the medical device manufacturer highlights a broader lesson for regulated industries: Generative Ai delivers meaningful returns when it is tightly coupled to existing governance structures and not when it attempts to replace them. What began as an experiment with a generic chatbot can evolve into a more deliberate configuration that uses team licenses with stronger security, richer prompts, and ultimately a RAG enhanced workflow that aligns better with regulatory expectations and internal risk tolerance.

By localizing models, curating domain-specific knowledge bases, and preserving human review at every critical step, organizations can reduce time to compliance, lower dependence on external consultants, and improve consistency in their regulatory narratives. For Fractional Technology Leaders, this pattern is a strategic lever because it creates a repeatable and defensible way to introduce Generative Ai into high stakes environments while demonstrating both technical fluency and regulatory empathy.

The implication is very clear: automation for its own sake is a poor fit for most regulated use cases, but targeted augmentation of human expertise is not. Leaders who frame Generative Ai as a disciplined extension of established processes, rather than a disruptive shortcut, will be better positioned to unlock value quickly, earn stakeholder trust, and build a portfolio of successes that can be replicated across clients and sectors.

Contacts:

For research inquires email :

research@galson.com

Follow Galson Research on social media: