PREPARING FOR MACHINE-ASSISTED REVIEW:

What Regulators' Use of Generative AI Means for Sponsors

SSI Strategy

Regulatory authorities around the world are actively integrating generative AI into their daily work. What was once speculative is now operational reality. In the United States, the FDA has rolled out Elsa, a secure AI assistant already in use across all centers. Elsa reflects the FDA's long-term AI/ML strategy to modernize regulatory review through secure, scalable, and internally governed tools. In Europe, REGULUS has emerged as a shared platform for over 20 EU and EEA agencies, offering a practical example of how AI can be integrated into regulatory workflows.

Regulators are already using AI and sponsors need to understand how far along those efforts are and what that means for submission readiness.

What Happens When Review Itself Is AI-Accelerated

Tools like Elsa and REGULUS are not theoretical, they are actively used by regulators. Reviewers now rely on AI to process high volumes of documentation, identify inconsistencies, and accelerate labor-intensive steps. The submission environment is shifting toward faster, more structured scrutiny, with increasing expectations for traceability, consistency, and quality.

FDA's Elsa and the Modernization of U.S. Regulatory Review

Elsa represents a leap in regulatory capability. It assists with reviewing clinical trial protocols, summarizing adverse event reports, automating label comparisons, generating code for nonclinical data handling, and prioritizing inspection planning.

Operating entirely within secure FDA infrastructure, it supports staff across centers, not just reviewers. Elsa's role in simplifying internal workflows signal the FDA's broader strategy: to modernize regulatory science through scalable, AI-enabled systems that uphold security and efficiency. FDA Commissioner Dr. Marty Makary has emphasized that tools like Elsa will shorten timelines for new therapy evaluations and enhance regulatory consistency.

By keeping sponsor data separate and focusing on internal optimization, Elsa balances innovation with trust. The takeaway for sponsors is clear: while submission formats like eCTD remain, the expectations for clarity, alignment, and analytical depth are increasing. Sponsors should also consider that Elsa's capabilities, such as inspection planning and faster label comparisons, can increase the intensity and precision of early scrutiny.

REGULUS and Europe's Shared Infrastructure for AI-Enabled Review

REGULUS is the first fully local generative AI platform deployed by a medicines agency. In the EU, REGULUS represents a region-wide step toward operationalizing generative AI in regulatory review. It is a secure, on-premise system built to comply with GDPR and the EU AI Act, capable of handling confidential regulatory material without cloud dependencies. REGULUS performs document summarization, translation, comparison, question answering, and template-based drafting, all under expert human oversight.

While REGULUS was originally developed by Sweden's MPA in under four months, it is now being adopted across the EU and EEA through the IncreaseNET initiative. It exemplifies how a shared capability can scale across borders while preserving national autonomy and legal safeguards. Rather than a standalone national tool, REGULUS has become a practical foundation for EU-wide regulatory AI use.

At the same time, the European Medicines Agency (EMA) is leading a multi-year effort to embed AI across the European Medicines Regulatory Network, guiding harmonized adoption and best practices across the EU.

From Strategy to Execution: How Other Agencies Are Moving

Beyond FDA and Sweden, other regulatory authorities are piloting or adapting LLM tools for focused tasks. EMA's roadmap focuses on training, ethics, and harmonization rather than central model development. Swissmedic, Japan's PMDA, and other agencies are engaged in pilots involving open-source or commercial LLMs. In China, national authorities are collaborating with state research bodies to build LLMs for compliance and legal review.

While full proprietary development remains rare, the global trend is toward responsible adaptation where regulators integrate powerful general-purpose models within domain-specific governance frameworks. There is growing consensus that generative AI will become a core part of regulatory infrastructure.

Lessons from Practice – Meridian One and the Role of Oversight

SSI Strategy's Meridian One initiative, a real-world application of generative AI for regulatory document generation, offers direct insight into implementation challenges and design principles.

"AI is not replacing regulatory professionals, but it is redefining their role," says Johan Strömquist, Chief Innovation Officer at SSI Strategy. "The professionals of the near future won't be drafters; they'll be workflow directors."

Meridian One was applied to high-value regulatory outputs, including Quality Overall Summaries, nonclinical overviews, and clinical efficacy sections. It was developed with strict guardrails: prompt discipline, content traceability, version control, and integrated human oversight. These elements were essential for generating consistent, review-ready materials.

"We learned that regulators aren't afraid of AI," Johan adds. "They want structure, transparency, and reliable oversight, the same standards that apply to traditional processes."

What Sponsors Need to Know Now

The structure of submissions isn't changing, tools like eCTD are still in place. But the way content is reviewed is evolving. AI tools enable regulators to detect subtle inconsistencies, unsupported logic, or fragmented narratives faster than before.

This increases the demand for clarity, internal alignment, and strategic cohesion across the dossier. While regulators accelerate, sponsors may face increased regulatory queries, mid-cycle adjustments, and greater resource demands after submission.

With different regulators moving at different speeds and in different ways, sponsors must stay alert to both official guidance and practical implementation. What counts as "AI-prepared" in one region may differ subtly from another. The real differentiator now is workflow readiness

Don’t Wait for a Mandate

Generative AI is already reshaping regulatory review. Sponsors should begin building internal AI capabilities, governance, training, validation frameworks, and tooling that support high-value use cases. Treat AI as a team member. Structure its role. Validate its outputs. Design for oversight. The submission environment has changed and it’s time to adapt.

Key Takeaways

- Regulators now use secure internal LLM systems to summarize, compare, and review submissions. Sponsors should assume their dossiers will be reviewed by both humans and machines.
- Automated tools expose narrative gaps and inconsistent cross-references that manual review might overlook. Sponsors need disciplined document architecture and traceable logic.
- The credibility of both regulator and sponsor AI use depends on version control, prompt discipline, and auditability.
- Sponsors who pilot governed AI use now will adapt more readily as machine-assisted review becomes standard practice.

The Authors

Ingela Loell
Scientific Writer & Editor

Johan Strömquist
President & Chief Innovation Officer

More Than a Service Provider

We are your dedicated partner in turning groundbreaking ideas into tangible solutions that make a difference. Combining our therapeutic expertise with strategic problem-solving, we help you overcome obstacles and reach every milestone on your journey.

Let's talk!

SSI Strategy

www.ssistrategy.com