SecureCHEK® AI
The Marketing Accelerator

TECHNOLOGY OVERVIEW · FULL HYBRID AI

Why a Full Hybrid Architecture Outperforms GenAI Alone

Accuracy, traceability, and governance for life sciences promotional and medical content

EXECUTIVE SUMMARY

Accuracy Requires More Than Generative AI

Generative AI is changing how life sciences companies create and review content, but on its own it cannot meet the accuracy and accountability that Medical, Legal, and Regulatory (MLR) review demands. Models deployed alone can produce plausible but incorrect statements, cite sources that do not exist, give different answers to the same question, and offer little insight into how a conclusion was reached.

Retrieval-augmented generation (RAG) reduces these risks by grounding responses in curated sources, but it does not eliminate them. For regulated content, accuracy requires a full hybrid architecture: generative AI combined with analytical AI, rules-based validation, retrieval grounded in a structured claims library, agentic workflows that harden into deterministic checks, and human oversight.

SecureCHEK AI is built this way by design. Each layer covers the weaknesses of the others, so every finding can be traced to its source and every decision remains with a subject matter expert.

Grounded
Every judgment is anchored to the client's own approved claims, references, and labeling.

Traceable
Every finding links back to the exact source passage, with a record of every decision.

Governed
Deterministic checks and humans in the loop keep final authority where it belongs.

THE CHALLENGE

The Limits of GenAI Alone

Published research consistently documents the same weaknesses when large language and vision-language models are deployed without additional safeguards:[1]

- **Hallucinations and factual errors:** plausible but incorrect statements, invented citations, and misstated facts.[2]
- **Stale or missing knowledge:** training cutoffs leave gaps, and outputs rarely show where information came from.

[1]Huang, L., et al. "A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions." ACM Transactions on Information Systems (2024). https://doi.org/10.1145/3703155

[2]Walters, W. H., and Wilder, E. I. "Fabrication and Errors in the Bibliographic Citations Generated by ChatGPT." Scientific Reports 13 (2023): 14045. https://doi.org/10.1038/s41598-023-41032-5

- **Inconsistency:** results vary with prompts, settings, and model versions, which undermines reproducibility.
- **Limited explainability:** it is difficult to see why a response was produced or to verify individual statements.
- **Bias and governance risk:** models can reproduce biases in their training data, and misuse is hard to trace.

GenAI-only tools are fast to deploy and flexible, but these failure modes are why grounding and verification are required wherever content must be accurate, substantiated, and defensible.

THE COMMON FIX

Why RAG Alone Is Not Enough

Many newer tools rely on RAG, which retrieves relevant documents in real time and supplies them to the model so its responses are grounded in curated data rather than training data alone. RAG helps, but its results depend on the quality and completeness of the underlying sources and on the retriever surfacing the right passage. If a source is outdated or a key passage is missed, the model can still produce an incorrect answer, and it can blend retrieved facts with invented detail.[3]

Structured knowledge layers, which organize facts into defined relationships, add another safeguard by letting outputs be checked against verified connections. They too must be carefully built and maintained. Both approaches reduce errors. Neither removes the need for strict verification and human review.

OUR APPROACH

Layers of Verification in a Full Hybrid Architecture

SecureCHEK AI combines complementary technologies so that no single method carries the full burden of accuracy:

Layer	What it does
Generative AI	Anthropic's Claude models, accessed through Amazon Bedrock, provide language understanding, reasoning, and a plain-language explanation of each finding. Anthropic's safety research and guardrails make these models well suited to regulated use.
Analytical AI	Detects page layout, groups related content, and extracts claims, tables, and images from complex documents in the major Latin-script languages of Europe and the Americas.
Rules-based validation	Enforces deterministic checks, including reference detection and management, ISI and safety language matching, and market-specific requirements.
Master Claims Library	The structured data layer. Every approved claim is linked to its supporting references and the exact substantiating passage, with metadata for product, market, and labeling changes.
Retrieval grounding	Grounds every generative AI judgment in the client's own approved claims, references, and labeling, rather than in the model's training data.

[3]Barnett, S., et al. "Seven Failure Points When Engineering a Retrieval Augmented Generation System." 3rd International Conference on AI Engineering: Software Engineering for AI (CAIN 2024). https://arxiv.org/abs/2401.05856

Layer	What it does
Agentic workflows	When a market is added, an AI agent reads its regulations and industry codes and builds a market profile. Requirements already covered switch on automatically, and new requirements are handled by a dedicated agent until they are hardened into permanent, auditable checks.
Human-in-the-loop oversight	Users can choose to accept content as written, mark it for changes, or capture it as a new claim or claim variant. Final authority stays with subject matter experts.
Governance and auditability	Full traceability of sources, versions, and decision logic, with a record of every disposition to support compliance and audit.

EXPLAINABLE RESULTS

Findings Users Can Act On

Every pre-check sorts findings into three categories, each presented with its supporting evidence so users can see exactly what was flagged and why:

Red Flags
Serious inconsistencies with approved product labeling.

Evaluate Critical
Medical or scientific inaccuracies that need correction or discussion.

Evaluate Other
Errors in non-claims content, such as editorial and formatting issues.

DEPLOYMENT

Built for Enterprise Environments

- **Flexible integration:** runs as a stand-alone platform or integrated with Veeva PromoMats, Aprimo, and Adobe InDesign.
- **Client-controlled hosting:** can be deployed in SecureCHEK AI's environment or in a client's own AWS account using the client's own model access, so content and claims libraries stay within the client's infrastructure.
- **Agent-ready:** a first release of our MCP connector exposes the core pre-check loop to AI agents, with access for additional platforms expanding as the connector matures.

CONCLUSION

The Most Defensible Design for MLR

Generative AI brings fluency and reasoning, but on its own it lacks the accuracy, consistency, and traceability that regulated content requires. RAG and structured knowledge reduce errors, yet both depend on careful curation and verification. A full hybrid architecture that combines generative AI with analytical extraction, rules-based validation, a structured claims library, agentic workflows that harden into deterministic checks, and human governance is the most defensible design for MLR today. It is the architecture SecureCHEK AI was built on.

AUTOMATE · ANNOTATE · ACCELERATE
securechek.ai