Leading MedTech Innovator

Generic AI failed them. Context-aware agentic AI cut investigations from 12 days to 3.

66%
FASTER COMPLAINT THROUGHPUT

96%
FIRST-PASS EXTRACTION ACCURACY

60%
LESS CASE RESEARCH TIME

<2%
PRE-SUBMISSION ERROR RATE

About the customer

A leading MedTech company developing innovative medical devices to improve patient outcomes. With growing regulatory demands and high complaint volumes, they needed to improve quality and complaint management without increasing headcount. They had tried generic AI — it delivered keyword-based search and surface-level summaries that nobody used for real decisions. They needed AI that understood regulatory intent, not just pattern matching.

Industry	MedTech / Insulin Delivery
Pain point	Generic AI, slow investigations

Challenge

The company had already tried AI. It didn't work — not because AI itself failed, but because **generic AI doesn't understand MedTech compliance**. Their previous tools delivered keyword-based search and surface-level summaries that provided little value in a compliance-critical environment. The quality team needed AI that understood regulatory frameworks, device-specific terminology, and investigation protocols. What they got was autocomplete.

The gaps were measurable and painful:

Key outcome	66% faster complaint throughput
Platform	Salesforce + Multi-LLM

Want more details? Download the full case study.

We'll walk you through Smarteeva with your use cases, your workflows, and your data.

Email address

Download case study ↓

- Lack of contextual understanding. Generic AI missed the meaning behind device-specific terms and regulatory language, leading to rework in complaints, CAPAs, and filings. Investigation cycles took 8–12 days versus the ideal 3–5 days. 15–20% of AI-assisted drafts required manual rework before submission.
- Limited automation that didn't integrate. Previous automation efforts didn't integrate deeply into the workflow, resulting in fragmented gains. Each complaint investigation consumed 6–8 hours of senior investigator time — and 40% of that was spent on data lookup and formatting, not actual analysis.
- Heavy manual data handling. Extracting structured data from unstructured sources — PDFs, emails, forms, medical records — remained labor-intensive and error-prone. Unit conversions, field standardization, and cross-market data harmonization required repeated manual intervention. Data preparation consumed 2–3 hours per complaint with a 5–8% error rate.
- Low business utility from existing AI. Previous AI tools had low adoption rates and were rarely used for high-stakes decisions. The team needed AI to understand regulatory intent — not just surface-level pattern matching. The AI was available, but nobody trusted it enough to use it.

> *"We didn't need another AI tool that generates summaries nobody reads. We needed an embedded co-worker that could handle the repetitive but complex investigations faster and more*

> *accurately than our senior investigators could alone."*

Approach

Smarteeva deployed agentic AI — not as an add-on, but as an embedded co-worker that runs inside the complaint management workflow. The fundamental difference: these AI agents understand device-specific terminology, regulatory frameworks (FDA 21 CFR Part 803, EU MDR/IVDR), and investigation protocols. They don't just process text. They understand context.

The deployment was designed around three principles that separated it from every previous AI attempt:

- Context-aware, not generic. Each AI session loads device history, prior complaints, and regulatory precedents into context before responding. The agent routes to the appropriate LLM model based on the task — OpenAI for language-heavy analysis, Gemini for multi-modal document processing. Follow-up investigations build on prior context, not restart from zero.
- Autonomous, not prompt-dependent. Agents run in the background — populating fields, values, and data without users typing a single prompt. The quality team doesn't operate the AI. The AI operates alongside them.
- Continuously learning from regulatory feedback. The agents track which investigation findings and compliance strategies pass FDA/EMA review without objection. Over time, prompt templates are refined based on regulatory acceptance patterns, and new case types are

incorporated into decision trees automatically.

Solution

Smarteeva deployed seven distinct AI capabilities — each embedded directly in the complaint workflow, each targeting a specific bottleneck that generic AI had failed to solve.

1. Context-aware AI with session memory. AI agents embedded in the CRM understand device-specific terminology and regulatory frameworks. Each session loads device history, prior complaints, and regulatory precedents. The agent maintains conversation history so follow-up investigations build on prior context — not restart from zero. Multi-LLM routing sends language-heavy tasks to OpenAI and multi-modal document processing to Gemini.

2. SmartExtraction — structured data from unstructured sources. Agentic AI extracts and standardizes fields from PDFs, medical records, complaint forms, and emails. Automatic unit conversion detects measurement units and converts to regulatory-standard formats. Field mapping translates customer terminology into standardized data models. Result: 96% first-pass accuracy on field extraction, reducing data validation from 2–3 hours to 15–20 minutes per complaint.

3. Similar-case intelligence via semantic matching. Instead of keyword search, Smarteeva's agents use semantic embeddings to surface regulatory precedents and similar investigation

findings — finding cases with similar device failures, patient outcomes, or investigation patterns even when terminology differs. Result: Investigators spend 60% less time on case research. Investigation cycle time reduced from 8–12 days to 3–5 days.

4 Dynamic compliance checklists with intelligent validation. AI agents auto-populate checklists based on device classification, incident type, and regulatory jurisdiction. Entries are validated against regulatory requirements in real time — ensuring all FDA Form 3500A fields are populated before submission. Missing information is flagged before it becomes a compliance gap. Result: Pre-submission error rate dropped from 12–15% to less than 2%. Time to regulatory-ready submission reduced by 40%.

5 Smart summaries for multiple cases. For products with recurring issues, agentic AI consolidates findings across 10+ complaints into a single pattern-analysis summary. Auto-generates template responses that investigators customize. Branches logic based on severity, device type, and regulatory pathway. Result: Multi-case analysis that previously took 3–4 days now completes in 2–3 hours.

6 Continuous refinement based on regulatory feedback. Agents track which findings pass FDA/EMA review without objection. Prompt templates are refined based on acceptance patterns. New case types are learned and incorporated into decision trees. Result: Resubmission rates for regulatory queries decreased by 35%.

7 Platform reliability at scale. Platform updates (API version upgrades) and defect fixes across checklists, similar

records, and decision-tree auto-completion increase stability as complaint volume grows. The system gets more reliable as it handles more — not less.

Results

The numbers tell the story. This wasn't an incremental improvement from adding AI features — it was a fundamental shift in how the quality team operates.

- 66% increase in complaint throughput — the team handles significantly more complaints without adding headcount.
- 96% first-pass accuracy on field extraction — data validation time reduced from 2–3 hours to 15–20 minutes per complaint.
- 60% less time on case research — semantic matching surfaces relevant precedents automatically. Investigation cycle time reduced from 8–12 days to 3–5 days.
- Pre-submission error rate dropped from 12–15% to less than 2% — intelligent validation catches gaps before they become compliance issues.
- Multi-case analysis from 3–4 days to 2–3 hours — pattern detection across 10+ complaints now happens in a single AI-generated summary.
- 35% reduction in regulatory resubmission rates — prompts refined based on what actually passes FDA/EMA review.

What's next

With agentic AI embedded across the complaint lifecycle and producing measurable results, the company is expanding Smarteeva's AI capabilities to:

- Adverse event trend analysis across the full product portfolio — detecting patterns that span individual complaint records
- Automated risk scoring for complaint severity classification at the point of intake
- Multi-market regulatory harmonization — using AI to align submission strategies across jurisdictions automatically
- Expanding Orchestra-built agents to new workflow areas including CAPA and risk management

Tried AI that didn't deliver?

See how context-aware agentic AI performs differently — embedded in the workflow, not bolted on. 30 minutes. Real data.

Book a Personalized Demo →