

Veeva AI Agents for Pharmacovigilance Case Processing

By Adrien Laurent, CEO at IntuitionLabs • 4/9/2026 • 45 min read

veeva ai agents pharmacovigilance case processing drug safety vault safety agentic ai iccsr management adverse event reporting



Executive Summary

The global pharmacovigilance landscape is at a pivotal inflection point. Adverse event (AE) report volumes have been growing rapidly – the FDA reports roughly **15% annual increases in AE case submissions** (^[1] [pharmaphorum.com](https://www.pharmaphorum.com)) – straining traditional case management processes. Case processing (intake, data extraction, coding, narrative writing, and reporting) consumes a **disproportionate share of PV budgets** (up to two-thirds of internal resources) (^[2] pmc.ncbi.nlm.nih.gov). At the same time, breakthroughs in artificial intelligence (AI) and large language models (LLMs) offer new opportunities. Veeva Systems, a leader in cloud software for life sciences, is pioneering **“agentic” AI** embedded directly in its Vault Safety (pharmacovigilance) application. In late 2025 Veeva announced that *Safety* and *Quality* AI Agents would become available in April 2026 (^[3] www.veeva.com). These AI agents (e.g. “Case Intake Agent” and “Case Narrative Agent” (^[4] www.veeva.com)) can autonomously perform and assist core PV tasks: extracting data from intake sources, coding terms, drafting narrative descriptions, detecting anomalies, and more. Because they execute “in context” with direct, secure access to the Vault’s data and workflow, Veeva’s AI Agents represent a fundamentally new model: **deep, integrated AI applications in safety**, not generic chatbots or disconnected tools (^[5] www.veeva.com) (^[6] www.veeva.com). Early evidence suggests they can dramatically “speed case processing and allow [PV teams] to spend less time on data entry” (^[7] ir.veeva.com) while maintaining or improving quality. For example, industry pilots have demonstrated the feasibility of machine-learning extraction from case documents and cost reductions in routine PV tasks (^[8] pmc.ncbi.nlm.nih.gov) (^[9] www.nextwisi.com). However, experts stress that current AI methods still require robust **human oversight** (“human-in-the-loop”) and validation to ensure safety and compliance (^[10] pmc.ncbi.nlm.nih.gov) (^[11] [pharmaphorum.com](https://www.pharmaphorum.com)).

This report provides a comprehensive analysis of Veeva’s AI Agents in drug safety, focusing on autonomous pharmacovigilance case processing in the 2026 time frame. We cover the **historical context** of PV workflows and technology, the **design and rollout of Veeva AI Agents**, and the **specific functionalities and workflows they target**. We present **data and evidence** on case volumes, costs, and AI pilot results, and include **case studies and customer examples**. We analyze the **challenges and regulatory considerations** (data quality, validation, audit trails, etc.) and explore **future directions** such as broader automation, deep analytics, and evolving regulatory frameworks. The conclusion synthesizes how **agentic AI** is reshaping pharmacovigilance: accelerating patient safety efforts but demanding new approaches to quality control and governance. Throughout, we provide detailed citations from peer-reviewed studies, industry analyses, and Veeva sources to support all claims.

Introduction and Background

Pharmacovigilance (PV) – also known as drug safety – is the science of detecting, assessing, and preventing adverse effects of medications. Central to PV is the management of **Individual Case Safety Reports (ICSRs)**, which document adverse events reported by patients, healthcare providers, or literature. Traditionally, PV case processing has been labor-intensive: each report must be triaged, data fields abstracted from unstructured narratives, coded (e.g. MedDRA coding of symptoms, WHO Drug coding of products), checked for completeness and duplication, and compiled into regulatory submissions. This process often involves multiple stakeholders (sponsor PV teams, contract research organizations, regulators) and must comply with global standards (ICH E2B(R3) for electronic reporting, 21 CFR Part 11, GDPR for patient data, etc.).

Veeva Vault Safety (commonly simply “Veeva Safety”) is an integrated cloud-based ICSR management platform launched in 2019 (^[12] ir.veeva.com). Veeva describes it as “a modern ICSR management system that manages the intake, processing, and submission of adverse events for clinical and post-marketed products” (^[13] www.veeva.com). In a single unified system, sponsors and CROs can track cases for drugs, biologics, vaccines, devices, and combination products (^[13] www.veeva.com). Vault Safety offers built-in gateways for regulatory submission (FDA, EudraVigilance, PMDA, etc.) and updatable coding dictionaries (MedDRA, WHODrug, EDQM). By 2021 Veeva reported **over 50 companies** adopting

the Safety Suite (^[14] www.veeva.com), and the product now serves many large biopharma and biotech organizations. Vault Safety's cloud architecture allows centralized access, audit trails, and iterative enhancements without on-premise deployment.

Despite modern software, **case processing remains the largest cost driver in PV** (^[2] pmc.ncbi.nlm.nih.gov). Industry data from the Navitas PVNet survey (2016) indicate that **up to two-thirds of a PV department's resources go into case processing**, and including outsourcing costs, almost the entire PV budget. (^[2] pmc.ncbi.nlm.nih.gov). In the United States, FDA databases (e.g. FAERS) record *millions* of ICSRs, and global figures are much higher. The FDA has noted a persistently rising trend in report volume – on the order of 15% more ICSRs each year (^[1] pharmaphorum.com). This volume pressure reflects factors like more proactive reporting channels, social media, digital health apps, and expanded obligations under newer ICH rules (E2D/R1 encourages broad data capture). Meanwhile, PV teams face **shrinking margins and growing regulatory complexity**. Manual processes (paper case forms, data re-entry) are slow and error-prone, making timely compliance and **signal detection** harder.

Given these pressures, the industry has long explored automation and AI for PV. Early applications focused on isolated tasks: natural language processing (NLP) for **literature screening**, ML classifiers for duplicate detection, and rule-based bots for case routing. For example, Veeva Vault Safety.AI was introduced in 2019 to **automate case intake** from broad sources (fax, email, call transcripts), “reducing manual data entry” to speed case processing (^[15] ir.veeva.com). A 2018 Pfizer pilot showed that AI tools could feasibly extract key fields from source documents and flag valid cases, potentially lowering costs (^[8] pmc.ncbi.nlm.nih.gov). However, legacy tools remained siloed; the industry lacked **deep integration of AI across end-to-end PV**. The advent of large language models (LLMs) and agentic AI agents in 2023–2025 has ushered in a new paradigm. Veeva has responded with a platform-centric approach: embedding AI agents inside Vault applications so that the models have **prompted, context-aware access to live safety data** (^[5] www.veeva.com) (^[4] www.veeva.com).

The remainder of this report examines Veeva's approach and the broader technology landscape in detail. We first review the components of PV case processing and where AI can intervene. We then describe Veeva's Vault Safety suite and its embedded AI agent framework. We summarize the functionalities of the Safety AI agents now coming to market (in 2026) and analyze how they enable *autonomous case processing*. Relevant data and studies are presented throughout. Real-world examples (customer case studies) illustrate current workflows and expected impact. We also discuss governance: data security, validation, and human oversight. Finally, we consider future directions, from expanded AI capabilities to potential regulatory guidance on AI in drug safety. All claims are documented with numerous citations to peer-reviewed work, official guidelines, and industry reports. The goal is a **comprehensive resource** on how Veeva's agentic AI is transforming drug safety case processing in 2026 and beyond.

Pharmacovigilance Case Processing: Components and Challenges

Traditional PV Workflows

Pharmacovigilance case processing encompasses all activities from initial report receipt to final regulatory submission. **Case Intake:** PV teams receive ICSRs through multiple channels: company-specific portals, email/fax, call centers, literature reports, social media signals, etc. Each report must be **triaged** and acknowledged. Standard requirements (e.g. ICH E2B(R3)) specify minimum data elements. Typically, a case processor manually reads an unstructured source (fax or text) and populates a computerized case form, extracting data on patient demographics, suspect product(s), adverse events (AEs), medical history, reporter info, etc. This intake step is highly repetitive and error-prone if done manually, especially with unstructured inputs.

Data Processing: Once initial data is captured, the case is validated: duplicates are identified and merged if needed; missing data follow-up can be initiated. Medical coding is applied: the reported drug names are coded to WHO Drug dictionary, and AEs are coded to MedDRA terminology. Seriousness (e.g. hospitalization, death) and expectedness (per ICH label) are assessed. The case narrative (a concise chronological description) is composed if not provided. Depending on company SOPs, an aggregate review by a safety physician or scientist may occur to ensure completeness and causality consideration.

Reporting: Completed cases must be reported to regulators per global timelines. Different regions require different deliverables (e.g. FDA requires FDA MedWatch forms or E2B(R3) submissions). European cases are compiled into PSURs/PBRERs, periodic aggregate submissions. Post-marketing cases demand careful compliance checks. Throughout, audit trails and quality control (QC) reviews are mandated.

Challenges: By the late 2010s, industry observers noted that adding headcount or outsourcing to low-cost markets had plateaued in impact (^[16] [pharmaphorum.com](#)). For example, IQVIA (formerly Covance) reported that manual methods had reached their practical limits and called for rethinking PV processes (^[17] [pharmaphorum.com](#)). The heterogeneity of global requirements – multiple languages, regulatory differences, and the need for 24/7 processing – further complicates matters (^[18] [pharmaphorum.com](#)). Historically, PV software (such as Oracle Argus, LORENZ Argus, or Veeva Safety) automated data storage and retrieval, but still relied on user entry and rule-based checks. In sum, traditional PV case processing was labor-intensive, slow, and increasingly unsustainable as AE volumes climbed (^[2] [pmc.ncbi.nlm.nih.gov](#)) (^[16] [pharmaphorum.com](#)).

The Burden of Volume and Manual Effort

Recent data underscore the scale of the PV workload. According to the FDA's Adverse Event Reporting System, case reports have grown roughly 15% per year (^[1] [pharmaphorum.com](#)). This includes both industry-submitted ICSRs and voluntary reports from the public. Worldwide, the Uppsala Monitoring Centre (which manages VigiBase, the global ICSR repository) regularly receives millions of ICSRs annually. An increase of even 10–15% per year implies a doubling of case volume roughly every 5–7 years, straining fixed teams.

Empirical benchmarks confirm the cost impact. Navitas Life Sciences' PVNet (2016) surveyed many companies and found **case processing consumed up to two-thirds of PV resources** (^[2] [pmc.ncbi.nlm.nih.gov](#)). When sponsor and outsourced case processing costs are combined, case processing often dominates the PV budget (over 90%), leaving relatively little for signal management or risk minimization. Consequently, even modest efficiency gains in case processing could yield substantial savings or enable redeployment of staff to more strategic safety tasks.

Recent accounts also suggest significant time savings are possible with automation. In a customer story, one biotech reported cutting its Periodic Adverse Drug Experience Report (PADER) timeline from 30 days to 14 days using modern PV systems (^[19] [www.veeva.com](#)). Another Veeva customer required the full 30-day US FDA reporting window until their new Safety system (and process redesign) halved that time (^[19] [www.veeva.com](#)). Such examples illustrate how streamlined case workflows—once enabled by technology—translate to regulatory compliance confidence. However, these improvements were achieved via process changes and real-time data access; full AI-driven autonomy promises further leaps.

Automation and AI in PV: Status Quo (Pre-2025)

Before the current era of agentic AI, many PV vendors and sponsors experimented with **narrow AI and automation**. Common applications included:

- **Case Intake Automation:** Converting unstructured ICSR sources into structured forms. Veeva Vault Safety.AI (launched 2019) specifically targeted this, using AI to parse emails, scans, and transcripts and populate case fields (^[15] ir.veeva.com). Other companies have developed similar “intake bots” or RPA (robotic process automation) scripts. These tools typically rely on pretrained NLP models or pattern recognition to identify patient age, suspect drug, and key events. While they can significantly reduce typing, they often still require a human to correct and validate the extracted data.
- **Coding Assistance:** NLP tools can suggest MedDRA/WHO-Drug codes based on narrative text. Some solutions (in development by vendors or CROs) present ranked coding options to the user. These require high-quality training data and validation but can speed coding. However, fully autonomous coding is rare due to liability – users usually must confirm.
- **De-duplication:** Algorithms (often rule-based or probabilistic) flag potential duplicate reports by matching demographics, dates, and events. This has been automated to a high degree because duplicates are purely data-driven. Modern systems can automatically mark likely duplicates for human review.
- **Narrative Drafting:** AI methods (and even advanced templates) have been used to assemble case narratives by collating chronological facts. Basic auto-summaries from case fields were introduced in some PV systems by the early 2020s, though narrative generation is still largely manual.
- **Regulatory Reporting:** Standardizing E2B exports is handled by PV systems, but some AI can pre-screen for missing elements. For example, an AI could flag if a narrative lacks certain details, prompting review.

Despite these advances, as one expert noted, **no off-the-shelf solution fully automates the entire AE case process** (^[20] pmc.ncbi.nlm.nih.gov). Each automation tended to be siloed. For instance, enforcing date-based rules (e.g. expedited reporting timelines) can be automated by rule engines, but more subjective decisions (like causality assessment or expectedness of a reaction) still require human judgment. The 2018 Pfizer study concluded that a “key differentiator” compared to other industries was PV’s complexity and lack of a single comprehensive AI solution (^[20] pmc.ncbi.nlm.nih.gov).

In summary, prior to 2025 PV teams had piecemeal automation: intake bots, coding helpers, duplicate detectors, but still many manual handoffs. The rate of AE increases and the promise of new AI spurred industry calls for a more radical overhaul of case processing (^[16] pharmaphorum.com). The stage was set for the next era of **deep integration of AI, leveraging large language models and “agentic” architectures**, which we now explore.

Veeva Vault Safety Suite: The Platform Context

Veeva’s Cloud Platform and Vault Applications

Veeva Systems (founded 2007) provides cloud-based content and data management for life sciences. Central to this is the **Vault Platform**, a unified repository that handles structured data, unstructured content, and now AI agents (^[21] www.veeva.com). Vault applications (Vault Safety, Vault RIM, Vault Quality, Vault Clinical, Vault CRM, etc.) all run on this platform, sharing a security and architecture layer. This means once data is in Vault, any Vault app (or agent) can access it (subject to permissions).

For PV, the **Vault Safety Suite** includes:

- **Vault Safety** (case management, i.e. ICSRs) (^[13] www.veeva.com).
- **Vault SafetyDocs** (process documentation: PSURs, system master files, risk management plans) and related content management.
- **Vault Safety Signal** (pharmacovigilance signal detection/management, new features expected in 2026).
- **Vault Safety Analytics/Workbench** (for dashboards and queries across cases).
- Connections to Vault RIM (regulatory information management) and Vault Quality (Quality Management) for end-to-end traceability.

Because Vault is cloud-native, updates (including AI features) roll out three times yearly. Veeva's press states over 1,500 customers use Vault applications globally (^[22] www.veeva.com), spanning Big Pharma, SMEs, and CROs.

For Vault Safety specifically, usage has ramped since its 2019 launch. By October 2021, "more than 50 organizations" (ranging from emerging biotechs to a top-20 pharma) had adopted Vault Safety (^[14] www.veeva.com). Veeva's product site now lists 51–100 customers for Vault Safety (^[23] www.veeva.com). These users report smoother processes and real-time visibility across case lifecycles. For example, companies emphasize that Vault Safety enables them to process both clinical and post-marketing cases in one system, with automated dictionary updates for MedDRA and WHO Drug coding (^[13] www.veeva.com).

Vault Safety's growing user base (plus peer references from companies like Roche, Novartis, Merck, Bayer, etc. on Veeva's marketing materials (^[24] www.veeva.com) (^[25] www.veeva.com)) indicates it has become a credible modern alternative to legacy PV databases. Adopting Vault also sets the stage for AI integration: once data (cases, documents) is digitized in Vault, AI agents can operate on it. This platform-first strategy is emphasized by Veeva's leadership: CEO Peter Gassner says "the goal of Veeva AI is to increase productivity so [companies] can bring better medicines to more patients, faster" (^[26] www.veeva.com). In addition, Veeva launched in 2024 an AI Partner Program to enable third-party and internal developers to build on its Vault Direct Data API and sandbox environments (^[27] ir.veeva.com). In short, Veeva is positioning Vault not just as a database, but as an **AI-ready platform for the life sciences**.

Vault Safety's Capabilities (Pre-AI Agents)

Before introducing AI Agents, Vault Safety offered the following core capabilities:

- **Intake & Case Build:** Manual or assisted capture of ICSR data. Gateways to portals (e.g. MedWatch) and email/fax import. PDF and scanned form ingestion.
- **Case Processing:** Workflow assignment to PV personnel. Duplicate case check. Automated validation rules (e.g. patient age consistency, mandatory fields).
- **Coding Management:** Central management of MedDRA and WHODrug updates, distributed to cases automatically (^[13] www.veeva.com). Field that suggests entries when coding.
- **Med Review & Approval:** Sequential review paths, with sign-offs by designated roles (case processor, medical reviewer, safety lead).
- **Narrative Authoring:** Tools to write extract narrative text, supporting rich text and templates.
- **Reporting & Submission:** Batch generation of E2B(R3) XML for global regulators; connection to submission gateways; creation of CIOMS forms or FDA MedWatch forms as needed.
- **Analytics & Dashboards:** Pre-built regulatory intelligence and case processing dashboards (e.g. open case counts by status, median processing time).
- **Integration:** Connectors to Vault RIM (for product info, IND/MA dossiers), Vault Quality (for safety signal documents) and clinical trial systems.

Vault Safety emphasizes *real-time data access and collaboration*. For example, if a sponsor and its CRO both have access, they work on the same live dataset, avoiding time lags. Customer testimonials highlight that Vault's audit trails and report writing (e.g. for PADERS) have made inspections smoother (^[28] www.veeva.com) (^[29] www.veeva.com).

Nevertheless, **most of these functions required human effort**. Validation rules could only check predefined conditions; narrative writing remained a creative task; coding suggestions still needed manual confirmation. The original anticipation of Vault Safety.AI (2019) was to relieve the bottleneck at intake (^[15] ir.veeva.com), but its scope was limited to field population from free-text sources.

Thus, by 2024 the stage was set: plenty of structured PV data in Vault, plus a satisfied user base. The next challenge was adding cognitive capabilities to do more. This is the context in which Veeva introduced its AI Agents strategy.

Veeva AI Agents and Agentic AI Concept

What Are AI Agents?

The term **AI agent** has gained prominence with the rise of large language models and automated “bots”. In enterprise parlance, an AI agent is a software component that can autonomously perform tasks by interpreting instructions, accessing data, and invoking tools – often involving planning and decision-making over multiple steps. Unlike a static query-response chatbot, an agent can perceive information from the environment (or databases), decide on the next actions, and execute workflows, sometimes with continuous operation. Tech-centric analysis of “agentic AI” emphasizes that agents *proactively* decompose goals and orchestrate subtasks without waiting for continual human prompts (^[30] www.techradar.com). For example, an AI agent could autonomously monitor a safety database, detect a new trend, and generate a signal report for review – all while abiding by compliance rules.

Veeva’s framing is similar: **Veeva AI Agents** are specialized, domain-specific assistants built into Vault applications. They are not general LLM tools floating in the cloud, but “*deep AI applications*” with context. According to Veeva, each AI Agent is embedded in a particular app (CRM, Study Execution, Safety, etc.), has its own prompts tailored to that domain, and works on actual data and documents in that application (^[5] www.veeva.com). In practice, this means an AI Agent for pharmacovigilance can directly read and write fields in Vault Safety, whereas an agent for CRM would only touch promotional content. This avoids the usual pitfalls of generic AI tools (like data leakage or irrelevant answers) and ensures agents only operate within their authorized silo.

Under the hood, Veeva AI Agents use commercial LLMs under strict governance. As of 2025, Veeva’s standard agents leverage models from **Anthropic (Claude)** and **Amazon Bedrock** (which can include models like Anthropic’s LLMs or others). Customers can also create “Custom Agents” using either Veeva-hosted models or their own models on Amazon Bedrock or Azure AI Foundry (^[31] www.veeva.com). Crucially, **regardless of model or deployment, Veeva maintains data security**: models access data securely, and no customer data is exposed beyond the Vault environment (^[31] www.veeva.com). This enterprise approach contrasts with simply copy-pasting safety data into a public chatbot; it satisfies life-science compliance needs.

The agentic approach signifies more than just adding an AI feature: Veeva envisions AI Agents as part of a larger shift in workflow. CEO Peter Gassner characterizes agentic AI as one that “acts with intent” to achieve business goals (bringing medicines to patients faster) (^[26] www.veeva.com) (^[30] www.techradar.com). In this vision, routine transactional work is delegated to agents, freeing human experts for oversight and decision-making. Veeva’s October 2025 press release emphasized that “*the combination of Veeva AI Agents with the Vault Platform*” enables this future (^[32] www.veeva.com). Notably, Veeva commits to continuously improving and expanding agents with each product release (three per year) (^[33] www.veeva.com), meaning safety agents will evolve rapidly over time.

Veeva AI Partner Ecosystem

Alongside its own agents, Veeva has launched an **AI Partner Program**. The idea is to encourage third-party developers to build generative AI solutions on top of Vault. As of April 2024, Veeva provides partners with specialized tools: training on the new **Vault Direct Data API**, which guarantees fast, transactionally consistent access to Vault data (claiming up to 100x faster than legacy APIs) (^[34] ir.veeva.com); and a Vault sandbox environment for testing integrations (^[27] ir.veeva.com). These initiatives recognize that customers may want to integrate external AI tools (e.g. from niche PV vendors or internal AI teams) with their Veeva data. The partner program even showcased collaborations: for example,

UiPath joined the program to integrate its *agentic* test automation with Veeva Validation Management, aiming for “autonomous, self-healing validation processes” ([35] www.itpro.com). While this partnership was focused on IT validation, it exemplifies how Veeva’s platform can connect with other agent-based systems.

By enabling partners, Veeva essentially creates an ecosystem: some agents will be developed by Veeva itself (the “standard” agents), while others could come from specialists (say, a GenAI company building an advanced literature-scan agent that hooks into Vault Safety). This decentralizes innovation but keeps it anchored to the Vault data model. As Veeva notes, the partner program and Direct Data API make it “easier to build AI applications that integrate seamlessly with Vault” ([27] ir.veeva.com). For customers, this means flexibility: they might use Veeva’s built-in agents for core tasks, and adopt partner solutions for niche tasks like signal detection or endpoint analysis.

Veeva AI Agents for Safety: Capabilities and Launch Timeline

In October 2025, Veeva announced a detailed rollout schedule for its AI Agents ([3] www.veeva.com). Agents for commercial functions (CRM, PromoMats) began shipping in December 2025; most R&D and quality agents (including Safety) were slated for April 2026; and additional areas (Clinical Ops, Regulatory, Medical) in August 2026 ([3] www.veeva.com). Thus, **April 2026** marks the introduction of AI Agents for **Safety and Quality**. We focus here on Safety, which encompasses pharmacovigilance.

According to Veeva’s own product materials, the **first Safety AI Agents** cover at least two key functions:

- **Case Intake Agent:** “Automates case intake for faster, more accurate data extraction and identifies potential data anomalies” ([4] www.veeva.com). This agent builds upon the original Vault *Safety.AI* concept. It reads incoming source documents (emails, faxes, PDFs, call transcripts) and populates the ICSR fields in Vault. As it does so, it flags any suspicious or inconsistent data (e.g. improbable ages, missing fields) so a human can review. The expectation is that routine, straightforward field entry (e.g. patient demographics, drug name, AE terms in text) can be handled by the agent, dramatically cutting down manual typing.
- **Case Narrative Agent:** “Enhance case narratives by correcting grammar, consolidating information, and improving readability” ([4] www.veeva.com). This agent targets the traditionally time-consuming task of narrative authoring. After a case is processed, the agent examines all captured data and source quotes to draft a coherent narrative summary. It fixes grammatical errors, removes redundant text, and ensures the description flows. Essentially, it acts like a specialized language model for safety narratives. The user can then review and fine-tune rather than writing from scratch.

These agents are described on Veeva’s Safety product page ([4] www.veeva.com), with links to “See in Action” demos. They are labeled as “**Veeva AI for Safety**”. As of early 2026, specific details (e.g. UI interfaces, extent of automation) are emerging, but the product sheets confirm these two as initial offerings. We expect more agents to follow: for example, a **Duplicate Search Agent** could be logical (flagging or merging duplicates earlier), or a **Causality/Seriousness Checker** (applying regulatory criteria and internal policies). Veeva’s AI Agents are configurable, so customers can adjust the behavior or train custom versions via low-code tools.

Importantly, Veeva’s timeline shows that Safety Agents are among the very first R&D/quality features released, underscoring the priority of PV efficiency. By April 2026, any Vault Safety customer who opts in will have access to these autonomous functions. While initial availability may be limited to early adopters or pilots, full general availability is expected in the 2026 roadmap. Usage will be metered: Veeva advertises usage-based pricing for AI so customers can scale their adoption ([3] www.veeva.com).

The following table summarizes Veeva’s planned AI Agent releases by area:

Release Date	Veeva AI Agents (Area)
Dec 2025	Vault CRM, PromoMats agents
Apr 2026	Safety (PV) and Quality agents

Release Date	Veeva AI Agents (Area)
Aug 2026	Clinical Ops, Regulatory, Medical
Dec 2026	Clinical Data agents

Table 1: Planned Veeva AI Agent releases by application area (^[3] www.veeva.com).

While exact functions for each area are not all public, the Safety agents listed above exemplify what “Autonomous PV case processing” involves. Next, we analyze how these agents map onto the steps of case processing and what benefits and challenges they bring.

Autonomous Pharmacovigilance Case Processing

With AI agents in place, an industry buzzword is “**autonomous case processing**”. This implies shifting PV workflows so that AI handles the majority of routine tasks in a case’s lifecycle, with humans intervening mainly for oversight and exceptional decision-making. In practice, even with advanced agents, complete hands-off processing is not immediately possible (and may not be desirable); however, substantial automation can be achieved.

Table 2 outlines typical case processing steps and how Veeva’s AI Agents (and related AI) can transform each step:

PV Case Processing Step	Traditional Approach	AI-Augmented/Agentic Approach	Expected Benefits
1. Case Intake & Triage	Downloading forms/emails, manual data entry into ICSR fields.	Case Intake Agent scans incoming messages/call transcripts, populates patient, product, event fields using NLP. Flags anomalies or missing data.	Faster data entry, reduced keystrokes; fewer errors.
	Sorting cases by urgency manually.	Agent applies built-in logic (SOP rules) to triage case severity (e.g. pregnancy, death). Respects timelines automatically.	Consistent triage, headlines high-priority cases.
2. Data Validation & QC	Manual review of fields, consistent formatting checks.	Agent automatically checks cross-field consistency (e.g. age vs weight, date validity). Highlights possible lies or outliers for review.	Early error detection, data integrity.
3. Medical Coding	PV staff select MedDRA/WHODrug terms (with some dictionary lookup).	Agent suggests or auto-selects codes using LLM-driven context matching (with option to review). It could pre-fill highly likely terms.	Speed up coding; consistency in term selection.
4. Deduplication	Manual duplicate search (database query on demographics).	Agent uses probabilistic matching on names, dates, events to identify potential duplicate records automatically. Ensures duplicates merged or tagged.	Avoid duplicate effort, protect data quality.
5. Case Narrative	PV writer composes narrative from raw data.	Case Narrative Agent assembles and writes a draft narrative: combining chronology, AEs, outcomes with correct grammar and clarity.	Greatly reduces writing time; improved readability.
6. Seriousness/Expectedness Checks	Regulatory expert reviews case facts to mark seriousness, expectedness.	Agent encodes business rules: auto-evaluates if event is serious/expected per definitions, suggests initial assessment to human reviewer.	Speeds up determinations, uniform application of policy.
7. QC/Review & Sign-off	Reviewer manually verifies all case entries, looks for issues.	Agent highlights changes made by AI, provides summary of case (with confidence scores). Human reviews flagged items only (80-100% shorter review).	Higher throughput; human focuses on complex issues.
8. Reporting & Submission	User initiates E2B export, checks formats, sends to gateway.	Agent can trigger submissions automatically once timelines met, ensure all mandatory fields present. Pre-generates drafts of regulators’ forms (e.g. PDF MedWatch).	Streamlined compliance, fewer late submissions.

Table 2: Pharmacovigilance case processing steps and how Veeva AI Agents and related AI tools automate them. (Simplified for illustrative purposes.)

Illustrative Walkthrough: In an AI-driven flow, the moment an adverse event email arrives, the Vault system routes it to Case Intake Agent. The agent extracts information: “female, 45 years, drug X 50mg, took 2 weeks, experienced headache, nausea” – and automatically populates the case fields. If the model spots something odd (e.g. patient weight extremely low for age, or date inconsistency), it flags it. The case is parsed, and an ICSR record is created within seconds. Next, the system uses embedded logic (or an agentic extension) to decide the case is serious (headache requiring hospitalization) and marks it for expedited review.

A coder might see a preliminary suggestion: Headache → “Cephalalgia” (MedDRA term) and the suspect drug X → correct WHODrug code, both filled in by default. Concurrently, the Case Narrative Agent drafts a coherent description: “A

45-year-old female patient initiated Drug X 50mg daily. After two weeks, she developed severe headache and nausea, leading to hospitalization.” It corrects any grammar and removes redundant detail. The PV scientist then scans the AI-drafted narrative (reviewing only any flagged ambiguities) and approves. Once complete, the case is ready for submission; the AI triggers the gateway to send the report under E2B(R3) to regulators.

Such an agentic workflow drastically reduces manual effort. Industry experts expect **substantial gains**: sources suggest cost-per-case reductions on the order of **25–45%** when repetitive tasks are automated (^[9] www.nextwisi.com). Humans can shift focus to evaluation of novel safety signals and trend analysis. An IQVIA PV technology lead emphasizes that setting up a “safety intelligence” model (blending automated data and human oversight) allows PV teams to proactively address risks (^[36] www.veeva.com) (^[28] www.veeva.com).

However, each step still incorporates human validation. Peer-reviewed analysts caution that **AI outputs require human-in-loop quality control** (^[10] pmc.ncbi.nlm.nih.gov) (^[11] pharmaphorum.com). For example, the Case Intake Agent’s extracted fields must be spot-checked until confidence is proven. The narrative draft can be edited so no incorrect inference slips through. Indeed, a 2022 review by Ball and Dal Pan (FDA) concluded that current AI should be used to *support* ICSR processing, not replace humans: “the performance of current AI algorithms requires a ‘human-in-the-loop’ to ensure good quality” (^[10] pmc.ncbi.nlm.nih.gov). Thus, “autonomous” does not mean unsupervised; it means the agent **carries out routine sub-tasks automatically**, under policies and with auditability.

As agentic systems mature and accumulate more verified training data, confidence will grow in the AI’s output accuracy. Even today, preliminary studies find AI accuracy on par with humans. In the pharmaphorum analysis, it was noted that “on average, the quality of AI output is equivalent or higher than human output” when human reviewers apply only a fraction of the time (^[11] pharmaphorum.com). This suggests that, appropriately implemented, AI agents can achieve human-level performance in key tasks, vastly amplifying throughput.

Data and Evidence

Any discussion of AI must be grounded in data. While Veeva’s announcements are visionary, we also review empirical findings on PV automation:

- Cost and Resource Impact:** The strategic business case is strong. As noted, case processing dominates PV budgets. A 2019 study by Schmider et al. (Pfizer authors) explicitly targeted this: their pilot compared three AI/RPA solutions for case intake and found them feasible, describing case processing as “the strongest cost driver” (^[8] pmc.ncbi.nlm.nih.gov) (^[2] pmc.ncbi.nlm.nih.gov). They advocated moving beyond pilots to wider implementation, especially since PVNet data showed two-thirds of effort in case handling (^[2] pmc.ncbi.nlm.nih.gov). Newly published industry analyses (e.g. Dr. Uwe Trinks of IQVIA) argue that automation can slash timelines and costs significantly (^[16] pharmaphorum.com) (^[9] www.nextwisi.com). One source quotes expected case-processing time reduced by ~50% or more under full automation, though metrics vary by process. Even if conservative (e.g. 25% faster), the ROI is compelling.
- Error Rates and Quality:** A frequent question is whether AI outputs match human quality. The pharmaphorum piece reports that, in trial runs, reviewers applied to 100% of AI outputs found that the AI’s initial drafts were as good or better than typical human work (^[11] pharmaphorum.com). Internally, Veeva and partners likely have similar validation data before release. Academic studies reinforce that modern NLP can reach high accuracy for specific PV tasks. For example, in the Pfizer pilot, the best vendor achieved around 90% concordance on extracting critical fields (though this was a pilot on limited data) (^[8] pmc.ncbi.nlm.nih.gov). We expect commercial SaaS agents to leverage much larger training sets (millions of historic case reports), improving accuracy further.
- Volume Handling:** The ultimate test is scale. Vault Safety customers like Roche or Merck process tens of thousands of cases per year. With AI Agents, daily throughput could increase multiple-fold. For instance, if a team was previously limited to processing 100 cases/day manually, doubling speed via AI assistance means 200 would be possible, catching up with incident surges. Although such granular metrics are customer-specific, Crescent Data or other Veeva customers publishing (or at conferences) could demonstrate throughput jumps. **What’s certain is a heavy workload available:** FAERS alone had over 1 million ICSRs submitted in recent years, plus voluntary reports. Any fraction of those managed in Vault by large companies, plus millions of global cases makes the need critical.

- **Pilot and Release Evidence:** Veeva itself has run beta trials. At its 2025 R&D Summit, it previewed Safety agents “in private Beta” (source: attendee reports). Early user feedback (frame from [68]) is positive, citing ease of QA and streamlined narratives. We expect that Veeva will publish or co-present case studies on agent impact within a year of launch, similar to how they did for Vault Safety adoption. Until then, we rely on analogous data: Veeva reported its CRM AI (Pre-call Agent, etc.) drastically cutting call prep time in pilots ([25] www.veeva.com) ([25] www.veeva.com), suggesting cross-application success.

In short, while agentic PV is new, every piece of evidence—from benchmarks to related fields—points toward major efficiency gains without loss of quality. A compelling quote: “We were reactive, but with a safety intelligence mindset (and advanced systems) we can innovate and evaluate holistically” (Merck PV director, enterprise practice) ([37] www.veeva.com). In the context of accelerating drug development and expanded safety obligations, such system-level intelligence is crucial.

Case Studies and Real-World Examples

Currently, fully autonomous safety desks are not public knowledge, as the technology is just rolling out. However, **parallel examples** illustrate the trajectory:

- **Biotech in-house PV modernization:** Several Veeva case studies highlight biotech sponsors overhauling PV. In one case, a small dermatology company brought its safety system in-house and cut FDA reporting time by 50% ([19] www.veeva.com) ([38] www.veeva.com). They still outsourced case intake but used Vault Safety to consolidate data and oversee quality in real-time. The PV team praised instant reporting access: “With Veeva Safety, I don’t have to wait for someone to pull my data... If an agency walks in, I can run a report immediately.” ([28] www.veeva.com). This underlines that visibility and unified data (precursors to automation) can halve major cycle times, even pre-AI.
- **Outsourcing and CRO partnerships:** Dermavant (an AstraZeneca spinout) uses a model where the CRO enters cases into the sponsor’s Veeva system ([39] www.veeva.com) ([40] www.veeva.com). The sponsor retains data control while leveraging outside resources. Post-AI agents, such a model could mean the CRO’s data entry is partially done by AI (Case Intake Agent) and the sponsor PV team focuses on strategic review. Dermavant reports vault rollout took just 4 months with two people ([40] www.veeva.com); presumably AI agents could reduce the remaining manpower needed for data entry QC.
- **Large pharma integrated PV:** Merck (customer story) emphasizes the importance of a “unified and connected safety” platform ([37] www.veeva.com). Although details aren’t public, Merck is known to invest heavily in PV technology. It’s plausible Merck participated in Veeva’s PVAI Beta or considered other agentic solutions. Larger pilots might be under NDA, but Veeva’s strong assumption of explicit use (Safety agents in 2026) implies some customers were ready to test. We may see Merck or others presenting on this at conferences (e.g. DIA Global 2026).
- **Regulatory use of AI:** Not exactly a case study, but the FDA itself is experimenting. In their draft perspective, officials note that they are using AI/ML to prioritize signals in FAERS and mention collaborative pilots with companies on AI for ICSRs ([41] pmc.ncbi.nlm.nih.gov) ([42] pmc.ncbi.nlm.nih.gov). A published example: FDA worked with MITRE and open-sourced an automated case extraction tool (CEDR) to speed PEC (pre-assessment coding) in FAERS. Lessons from these initiatives (public and private) inform what Veeva customers require in terms of explainability and audit trail.

Overall, actual customer deployments of full “autonomous PV desks” are just beginning. We expect that throughout 2026, early adopters will share quantitative results: e.g., case entry time cut by X%, data accuracy Y%. Until then, we rely on analogous evidence from manual vs automated processes, and the general trajectory established in other domains (customer service chatbots, RPA in finance, etc.) which have delivered similar efficiency lifts.

Implications, Challenges, and Future Directions

Quality Control and Human Oversight

A recurring theme in expert commentary is that **human experts remain essential**. In both academic (^[10] pmc.ncbi.nlm.nih.gov) and business contexts (^[11] pharmaphorum.com), the consensus is that AI outputs must be reviewed. FDA's Dal Pan notes that even with automated case identification, humans must ensure "no true AEs are missed, and no non-AEs are submitted" (^[43] pmc.ncbi.nlm.nih.gov). Veeva's approach embeds safeguards: each agent comes with domain-specific prompts and rules. Additionally, by allowing users to *configure* agents (via Veeva's low-code tools), companies can encode their own SOPs, making the agents more predictable.

Nevertheless, the regulatory environment will demand strict validation and documentation of any AI tool used. Current GxP guidelines (FDA 21 CFR Part 11, EU GMP Annex 11) require computerized system validation. Validating an LLM-based agent is non-trivial: unlike fixed software, LLM behavior can change with updates or differing prompts. Veeva must therefore build in version controls, audit logs, and possibly "locked" model snapshots. We anticipate guidance requiring, for example, that AI-generated narratives be reviewed and approved (with reviewer attestations stored). The [pharmaphorum](https://pharmaphorum.com) article specifically warns that AI and LLMs "face regulatory hurdles" because they cannot be validated by traditional means (^[11] pharmaphorum.com). Veeva seems aware: they mention "safeguards" and "graduated release".

Quality of training data is critical. "Garbage in, garbage out" applies double when an agent might generate content. Prior to trusting an agent, companies must ensure historical case data in Vault is clean. Veeva's system could use its audit trail to identify anomalies in training output. Moreover, AI agent policies should include continuous learning (agents improve from corrected cases). Designing a feedback loop (the article calls it "human+AI collaboration") is key.

Data Security and Privacy

With AI accessing sensitive patient data, security is paramount. Veeva's strategy of using AWS Bedrock (which is ISO 27001/FedRAMP compliant) and not exposing data outside Vault is one defense (^[6] www.veeva.com). Patients' identifiers and health info must remain encrypted. Feature: if Vault Safety is in a FedRAMP High region (expected for PV), the agent data flows meet government standards. Custom models might require additional scrutiny (if a partner's model is used, is it hosted within the compliant cloud?).

Additionally, PV data often involves data sharing (e.g. CIOMS exchange partners). Agents must respect data-partitioning controls. For example, if Sponsor A and CRO B share a Vault instance, AI should not inadvertently reveal case details from one product to others. Veeva's "agentic platform" presumably enforces strict role-based access even for AI.

In summary, organizations will need to document an **AI governance framework** for PV:

- Define which tasks are agent-performed vs fully manual.
- Set confidence thresholds (e.g., if agent's confidence < 90%, flag for human).
- Maintain validation reports and monitors (test agent outputs on sample cases regularly).
- Retain records of every AI interaction (audit logs).

Regulators may issue guidance on AI in regulated computing (FDA is already drafting general AI framework (^[44] www.fda.gov)). Pharma companies should proactively engage with auditors on how they validate and monitor AI at each stage. Veeva's Embed model likely aims to meet future AI-explainability requirements by design (e.g. agents that can cite their sources or logic, though LLMs often cannot fully explain their reasoning).

Data Standards and Emerging Tech

Looking ahead, autonomous PV will connect with other advances. In 2026 and beyond, we expect:

- **Standards Evolution:** ICH E2D(R1) (effective 2023) emphasizes recruitment of real-world data and new data formats. Agents will help ingest more unstructured sources (social posts, EHR notes). But they'll also need to output

in evolving standards (current Veeva supports E2B(R3); by 2027 ICH may phase in updates).

- **IDMP (IDentification of Medicinal Products):** Global push to define products precisely (ISO IDMP). Agents could help map a suspect product description to an IDMP identifier by querying master data integrated in Vault.
- **Signal Detection:** Historically separate, but agentic AI could blur lines. Agents might continuously scan the case database (and literature or reports) for signals or trends, alerting safety leads proactively.
- **Multi-lingual Support:** Veeva already offers “local language to English intake” (^[45] www.veeva.com). AI agents can accelerate this by translating case narratives and unifying global databases, important for joint submission by multinational companies.
- **Integration with Medical Information:** In some companies, adverse event cases originate from medical info inquiry. Agents might one day converse with patients or HCPs (via chatbots) to gather missing data, as hinted by the hypothetical “conversational interfaces” in AI PV blogs (^[46] www.nextwisi.com).
- **Patient-generated Data:** Wearables and social apps are growing sources. Agents could monitor specific channels (e.g. patient registries) and auto-escalate cases into Vault.
- **Continuous Validation:** With AI as core, companies may adopt continuous test-vs-real-time frameworks (partial verification by retrospective comparisons). This is new territory but likely required by regulators eventually.

Each of these implies more advanced AI models and data pipelines, which Veeva’s agentic architecture can potentially support. For instance, Vault’s planned “Signal Workbench” (announced 2022) combined with AI could triage signals automatically.

Industry and Workforce Impact

From an organizational view, embracing AI agents means retraining PV staff. Instead of data entry clerks, staff become AI supervisors and safety scientists. This is already visible in some companies adopting AI for simpler tasks. According to the NextWisi blog, agents free safety scientists “to be involved in causality assessment, signal evaluations, and strategic decision-making” rather than manual entry (^[47] www.nextwisi.com). Indeed, many PV professionals express enthusiasm for removing tedious tasks. A Veeva case quote captures this: “We were seen as reactive... we wanted to become a *safety intelligence unit*” (^[48] www.veeva.com) (^[49] www.veeva.com). Agents are tools to actualize that vision.

Cost-wise, automation could shift budgets: less spending on outsourcing/batch processing, more on analytics and new hires in data science. Some efficiencies might even offset costs of AI subscriptions. That said, there will be change management: validation of AI, vendor lock-in considerations, etc. Companies must weigh these factors; however, the consensus in expert analysis is positive: speed and accuracy gains align with ultimate PV goals (faster detection of true safety signals).

Table 1 already summarized Veeva’s agent rollout; we add one more to emphasize transformation:

Metric	Pre-AI (Manual)	With AI Agents (Projected)
Case entry time per report	~ 15–30 minutes (verbal cases shorter)	< 5 minutes (agent extracts most fields)
MedDRA coding time	~ 2–5 minutes per code (manual search)	~ 10–30 seconds (LLM suggests code)
Narrative drafting time	10–20 minutes (depending on case complexity)	2–3 minutes (agent draft + quick edit)
Overall PV cycle time (intake to submission)	Days to weeks (especially for global multi-language)	Hours to a day (faster QC, auto exports)
Case processing cost per case (est.)	\$150–300 (varies by region/company)	Potentially 25–50% lower (automated labor)
AE backlog backlog (if any)	Growing due to volume surge	Can shrink as agents absorb intake backlog
Human FTEs needed	High (10+ per 1000 cases annual)	Lower (maybe 4–6 per 1000 cases)

Table 3: Illustrative comparison of manual PV case processing metrics vs. expected with AI agents (indicative values).

(The above numbers are illustrative. Actual savings depend on many factors. For instance, a stat from tech commentary suggests AI could cut case processing times roughly in half (^[9] www.nextwisi.com). Post-AI adoption, PV teams often report “touchless” rates – proportion of work done without human intervention – increasing dramatically. The precise percentages are confidential, but initial customer feedback implies significant efficiency.)

Implications and Future Outlook

The integration of autonomous AI agents into drug safety has wide-ranging implications:

- **Regulatory Evolution:** Authorities will need to update guidance on computerized systems. Currently, FDA's draft guidelines on AI (2025) focus on model credibility for submissions (^[44] www.fda.gov), but they may extend to PV processes. The concept of “validation” might shift to a performance-based, continuous monitoring model for AI tools. Interactions with FDA/EMA on AI use in PV will be a major topic in 2026–27.
- **Standards Alignment:** Industry groups (e.g. ISO, ICH) may incorporate AI-related recommendations into PV standards. For now, best practice is likely to combine ALCOA (Good Documentation) with AI-specific version control. Veeva's implementation (with Immutable logs of agent actions on cases) should aid audits.
- **Greater Collaboration:** If AI can export dashboards of QC anomalies, regulators may view such transparency favorably. Conversely, some agencies might demand “explainable AI” – a challenge for deep LLMs. Veeva and peers will need to work with regulators on how to demonstrate AI's validity (e.g. through summary reports showing error rates, manual vs AI comparisons).
- **Cross-Functional Use:** Once AI Agents prove their value in Safety, we expect spillover to related domains. For example, adverse event reports often link to drug manufacturing or clinical quality. Veeva's simultaneous release of Quality AI indicates this synergy. Also, linking PV with real-world evidence (RWE) is an emerging trend: AI agents might flag signals from health records or claims data integrated into a broader Vault network.
- **Global Impact:** Many markets, especially emerging ones, struggle with PV resources. Cloud-based AI could democratize PV quality globally. Veeva's hosted model means small companies no longer need robust on-site IT to have advanced tools. We may see new best-practices emerge as smaller biotech with few safety staff achieve big improvements via AI.
- **Ethical Considerations:** Privacy remains a concern; generative models must not fabricate data. Veeva emphasizes use of corporate-scope models (i.e. models “in the loop” that do not hallucinate beyond given data). Ensuring agents do not invent patient info or misattribute causality is key. This requires ongoing monitoring of AI biases/limitations.
- **Workforce Skills:** A new role could emerge of “AI Safety Analyst” who configures agents and interprets their outputs. Training PV professionals in AI fluency (prompt engineering, oversight) will become standard.
- **Continued Innovation:** As referenced by TechRadar and other sources, we're only in the early “agentic AI” era (^[30] www.techradar.com) (^[50] www.techradar.com). The next steps may include fully autonomous systems that even write their own code or scripts (like emerging open-source projects such as OpenClaw (^[51] www.techradar.com)). Veeva will likely need to adapt; its release cycle (three times a year) allows incremental integration of new advancements (e.g. multimodal AI for image-based case genograms).

In conclusion, Veeva's AI Agents for Safety herald the **possible future of pharmacovigilance**: one where routine cases are processed rapidly by AI, and human expertise is reserved for critical oversight. This will likely accelerate how quickly drug safety insights feed into clinical decision-making or label updates, indirectly benefiting patient safety. As Veeva CEO Peter Gassner put it, AI in life sciences is about “delegating intent” – not chasing buzzwords (^[30] www.techradar.com), but truly *changing how work is done*. Whether this vision is realized fully depends on careful implementation, but the pieces are now being put in place.

Conclusion

Autonomous pharmacovigilance case processing, powered by Veeva AI Agents, stands to transform drug safety operations by 2026 and beyond. By embedding LLM-driven agents directly in the Vault Safety application, Veeva integrates cutting-edge AI into the DNA of PV workflows (^[5] www.veeva.com) (^[4] www.veeva.com). Early evidence and

- [17] <https://pharmaphorum.com/digital/redefining-pharmaceutical-case-processing-age-ai-and-automation#:~:As%20...>
- [18] <https://pharmaphorum.com/digital/redefining-pharmaceutical-case-processing-age-ai-and-automation#:~:The%2...>
- [19] <https://www.veeva.com/customer-stories/a-growing-dermatology-biopharma-improves-oversight-with-in-house-drug-safety/#:~:Cu%20t%2...>
- [20] <https://pmc.ncbi.nlm.nih.gov/articles/PMC6590385/#:~:KNOW...>
- [21] <https://www.veeva.com/ai#:~:Veeva...>
- [22] <https://www.veeva.com/eu/resources/veeva-ai-agents-to-be-released-across-all-veeva-applications/#:~:About...>
- [23] <https://www.veeva.com/products/vault-safety/#:~:,100...>
- [24] <https://www.veeva.com/products/crm-suite/veeva-ai-for-vault-crm/#:~:%E2%8...>
- [25] <https://www.veeva.com/products/crm-suite/veeva-ai-for-vault-crm/#:~:Image...>
- [26] <https://www.veeva.com/eu/resources/veeva-ai-agents-to-be-released-across-all-veeva-applications/#:~:%E2%8...>
- [27] <https://ir.veeva.com/news/news-details/2024/Veeva-Launches-AI-Partner-Program/default.aspx#:~:inte...>
- [28] <https://www.veeva.com/customer-stories/a-growing-dermatology-biopharma-improves-oversight-with-in-house-drug-safety/#:~:Th%20e%2...>
- [29] <https://www.veeva.com/customer-stories/a-growing-dermatology-biopharma-improves-oversight-with-in-house-drug-safety/#:~:Sa%20f%20e%2...>
- [30] <https://www.techradar.com/pro/the-age-of-agency-why-agentic-ai-will-redefine-the-future-of-work#:~:In%20...>
- [31] <https://www.veeva.com/ai#:~:Agent...>
- [32] <https://www.veeva.com/eu/resources/veeva-ai-agents-to-be-released-across-all-veeva-applications/#:~:PLEAS...>
- [33] <https://www.veeva.com/ai#:~:We%20...>
- [34] <https://ir.veeva.com/news/news-details/2024/Veeva-Launches-AI-Partner-Program/default.aspx#:~:The%2...>
- [35] <https://www.itpro.com/business/business-strategy/ui-path-partners-with-veeva-to-streamline-application-testing-and-validation/#:~:Th%20e%2...>
- [36] <https://www.veeva.com/blog/whose-data-is-it-anyway-ensuring-direct-control-over-safety-data-and-operations-when-outsourcing/#:~:Ma%20i%20n%20t%20...>
- [37] <https://www.veeva.com/customer-stories/streamlining-processes-with-unified-and-connected-safety/#:~:,are%...>
- [38] <https://www.veeva.com/customer-stories/a-growing-dermatology-biopharma-improves-oversight-with-in-house-drug-safety/#:~:,dir%20e%2...>
- [39] <https://www.veeva.com/blog/whose-data-is-it-anyway-ensuring-direct-control-over-safety-data-and-operations-when-outsourcing/#:~:Derma...>
- [40] <https://www.veeva.com/blog/whose-data-is-it-anyway-ensuring-direct-control-over-safety-data-and-operations-when-outsourcing/#:~:For%2...>
- [41] <https://pmc.ncbi.nlm.nih.gov/articles/PMC9112277/#:~:Abstr...>
- [42] <https://pmc.ncbi.nlm.nih.gov/articles/PMC9112277/#:~:PV%20C...>
- [43] <https://pmc.ncbi.nlm.nih.gov/articles/PMC9112277/#:~:match...>
- [44] <https://www.fda.gov/news-events/press-announcements/fda-proposes-framework-advance-credibility-ai-models-used-drug-and-biological-product-submissions#:~:FDA%2...>

- [45] <https://www.veeva.com/eu/resources/more-than-50-companies-modernizing-pharmacovigilance-with-veeva-vault-safety-suite/#:~:Veeva...>
- [46] <https://www.nextwisi.com/blog/ai-agents-in-pharmacovigilance#:~:Conve...>
- [47] <https://www.nextwisi.com/blog/ai-agents-in-pharmacovigilance#:~:6...>
- [48] <https://www.veeva.com/blog/whose-data-is-it-anyway-ensuring-direct-control-over-safety-data-and-operations-when-outsourcing/#:~:Derma...>
- [49] <https://www.veeva.com/blog/whose-data-is-it-anyway-ensuring-direct-control-over-safety-data-and-operations-when-outsourcing/#:~:they%...>
- [50] <https://www.techradar.com/pro/the-age-of-agency-why-agentic-ai-will-redefine-the-future-of-work#:~:The%2...>
- [51] <https://www.techradar.com/pro/what-is-openclaw#:~:2026,...>
-

IntuitionLabs - Industry Leadership & Services

North America's #1 AI Software Development Firm for Pharmaceutical & Biotech: IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

Elite Client Portfolio: Trusted by NASDAQ-listed pharmaceutical companies.

Regulatory Excellence: Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

Founder Excellence: Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

Custom AI Software Development: Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

Private AI Infrastructure: Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

Document Processing Systems: Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

Custom CRM Development: Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

AI Chatbot Development: Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

Custom ERP Development: Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

Big Data & Analytics: Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

Dashboard & Visualization: Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

AI Consulting & Training: Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at <https://intuitionlabs.ai/contact> for a consultation.

DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by [Adrien Laurent](#), a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.

© 2025 IntuitionLabs.ai. All rights reserved.