

CASE STUDY

# Graph AI

Rebuilding one of pharma's most manual workflows with GenAI-native SaaS

## Discover:

- How Graph AI's customers achieved >90% efficiency gains
- Why you should align your pricing roadmap with your automation roadmap
- Graph AI's two-and-a-half-week proof of concept

**Bessemer Venture Partners**

Copyright © 2026



# **Graph AI: A service firm turned AI-native solution for pharma and life sciences**

|   |          |
|---|----------|
| <b>Graph AI's pivot from AI services to GenAI-native SaaS</b> | <b>3</b> |
| <b>Key founder lessons</b>                                    | <b>4</b> |
| <b>The challenge</b>  | <b>5</b> |
| The hidden cost of keeping drugs safe                         |          |
| Graph AI's two-and-a-half-week proof of concept               | <b>6</b> |
| <b>The solution</b>   | <b>6</b> |
| Building for a domain that doesn't forgive mistakes           |          |
| A pricing roadmap that follows the level of automation        | <b>8</b> |
| <b>The result</b>   |          |
| Graph AI's three routes to market                             | <b>9</b> |
| What drug safety reveals about where AI wins                  |          |



Every drug on the market comes with ongoing regulatory reporting obligations. Once a pharmaceutical company wins FDA approval and begins selling treatment, the work doesn't stop; it multiplies. Patients around the world report side effects, clinical trials surface new complications, journals publish new findings, and regulators across countries issue new requirements. All of which has to be collected, classified, reviewed, and reported continuously by a team of people, for as long as the drug is on the market. This process, known as pharmacovigilance—or drug safety monitoring—is one of the most consequential and least-visible functions in the global pharmaceutical industry.

Yet, for most pharma companies, this critical function is outsourced to legacy IT services players and runs on a patchwork of different software tools, multiple layers of review, and armies of Pharmacology graduates doing loads of data entry. In many organizations, the real workhorses are Word documents and Excel spreadsheets, with a single adverse event taking over three hours due to the manual process. This is the world Graph AI's founders entered. Four technologists with no background in pharmacology and life sciences saw a new process to solve a legacy industry problem. While the AI solutions built were trained on domain data and precision-tuned for this exact regulatory context, Graph AI's new workflow and system—reimagined by SaaS veterans—allowed users to do in a minute what takes a team of dozens half a day.

**In our case study series, Launching AI products that win, we profiled Graph AI, which is building the first AI-native Patient Safety Operating System, purpose-built for pharmaceutical, biotech, and life sciences companies navigating the full adverse event lifecycle. We spoke with CEO Raghavendra Parvataraju and CTO Vijay Ponukumati, who walked us through their product story, the domain challenge they encountered and tackled, and how they transitioned from services to an AI-native SaaS platform.**

## Graph AI's pivot from AI services to GenAI-native SaaS

|                  |  |
|------------------|--|
| <b>Situation</b> | Graph AI was founded in 2024 by four enterprise technology veterans: Raghav Parvataraju (CEO, GSI background at LTI Mindtree, Infosys, Wipro), Vijay Ponukumati (CTO, ex-Google engineering lead for Google Cloud networking), Mohan Konyala (CPO, ex-ServiceNow), and Ashutosh Bordekar (CFO, ex-LTI Mindtree and Mastek). They initially launched as a services-only firm with a thesis around helping large enterprises navigate their AI journeys, anchored by their GraphX agentic AI platform.   |
| <b>Challenge</b> | Their first customer, a life sciences company in pharma, spent 18 months with IT services firms and failed to deliver AI workflows for patient safety, citing the regulatory complexity and precision requirements LLMs couldn't meet out of the box. The incumbent software players (Oracle, Veeva Systems, ArisGlobal) were only automating the software layer, not the human labor. The true incumbents were Microsoft Word, Excel, and people.   |
| <b>Solution</b>  | Graph AI built an MVP in two and a half weeks. They went deep on domain expertise, studying pharmacology, onboarding a medical doctor, and fine-tuning Gemini models on knowledge graphs purpose-built for the life sciences domain. They pivoted from services to a gen AI-native SaaS platform, Graph Safety, that unifies 14 functional process areas into a single unified platform with explainability, audit trails, and near-perfect accuracy baked in.   |
| <b>Result</b>    | Enterprise customers on the platform are now monitoring over 300 drugs, with efficiency gains of 90%+ versus the status quo. Total Cost Savings were 65% from the current spend using Graph Safety. Processing time per event went from over three hours to 10 minutes with humans in the loop. Google saw significant value that Graph Safety delivered on the GCP platform and how Google Vertex AI and Gemini models were fine-tuned for the life sciences industry and onboarded them as a partner on Google's Cloud Platform, with Graph Safety featured on Google Cloud Marketplace. |



## Key founder lessons from Graph AI's story

### 1. When no one has solved the domain problem, that's the domain opportunity.

The patient safety space was unsolved because it required deep domain expertise that other tech companies weren't willing or able to develop. Graph AI made that investment early (hiring an MD, studying pharmacology workflows from scratch), and it ultimately became their defensible moat.

### 2. A failed incumbent implementation can be the best possible lead.

Graph AI's first customer came to them after an IT services firm spent 18 months trying to deliver and couldn't. That failure created urgency, trust, and a fast pilot cycle: two and a half weeks to a working MVP. Early adopters with recent bad experiences are often the most motivated (and ready) to validate a new approach quickly.

### 3. Let your pricing roadmap follow your automation roadmap.

When Graph AI entered the pharmacovigilance market, the industry was pricing based on headcount. Existing software vendors charged only for what their software did, but the human labor on top was a separate cost entirely. Graph AI's insight was that pricing should expand in lockstep with how much of that human layer the platform automates. The result is a pricing architecture that grows naturally with product depth, gives customers a clear and predictable cost structure at every stage, and keeps Graph AI's revenue expansion tied to the value they're creating—not arbitrary seat counts or legacy SaaS benchmarks.

### 4. In a regulated industry, compliance is a product decision.

Complying with the FDA's CFR 21 Part 11, achieving SOC 2, and passing a computer system validation process is table stakes in pharma—but it's also a built-in barrier to entry that most AI startups won't clear. Graph AI treated the four-month validation process as product work, not overhead, and built auditability and provenance natively into the platform from the start. In industries where regulators effectively define the product spec, the companies that embrace those constraints earliest will be the hardest to displace.



## The hidden cost of keeping drugs safe

Pharmacovigilance is a regulatory mandate where every pharma company, biotech, and clinical research organization (CRO) operating in global markets is required to track adverse events and report them to regulators like the FDA and European Medicines Agency. Miss a report, and the fines have reached upwards of \$140M. Miss enough of them, and the drug comes off the shelf.

Today, delivering pharmacovigilance requires stitching together 40 to 50 different software tools across 14 functional process areas. The dominant platforms in most organizations aren't Oracle or Veeva. As Raghav puts it, "A lot of incumbency isn't from Oracle, Veeva, or Aris Global. It's from Word documents, Excel sheets, and PowerPoints. Those are the real incumbents in the space"—held together by human judgment and effort.

Not only are the tools specific and manual, but the workflow itself also runs through several distinct personas:

- Data entry professionals take incoming reports—from call centers, emails, clinical trial data, partner feeds, public databases, and regulatory XML feeds—and enter them into safety databases
- Pharmacology graduates triage each event, determining whether it qualifies as an adverse event at all
- Quality reviewers, typically with master's-level pharmacology training, audit that work
- And finally, medical reviewers or physicians sign off before anything goes to a regulator

Every handoff is manual, and every step is sequential.

### The data problem that compounds everything

Adverse events don't arrive in clean, structured formats. They come in as unstructured text from patient emails, journal articles, clinical study reports, and reports from healthcare providers. These distinct sources feed into the system, each with its own format, language, and level of completeness. Volume grows as pharma companies launch more drugs in more markets. And the workforce handling it often lacks the deep domain expertise the work demands. Many of the people doing front-line classification in call centers are junior pharmacology graduates making judgment calls they're not fully equipped to make.

The existing software players weren't helping solve this problem. They had built tools to manage the software layer of the workflow, such as case management, database entry, and submission formatting. But they weren't touching the human labor layer of automating the triage, the data extraction, the quality review. They were, in effect, digitizing the paperwork while leaving pharma companies to still pay for headcount, compliance overhead, and operational inefficiency. For pharma companies trying to move faster, launch more products, and respond to a regulatory environment that had grown significantly more complex since COVID, the status quo wasn't just inefficient. It was a ceiling.



## **Graph AI's two-and-a-half-week proof of concept**

Graph AI didn't initially set out to build a drug safety platform. When Raghav and the team founded the company, their thesis was straightforward: large enterprises needed help navigating their AI journeys. A team with deep expertise in enterprise technology—backgrounds from Google, ServiceNow, and major global systems integrators—was well-positioned to guide them. Their vehicle was GraphX, an agentic AI platform for building enterprise workflows, and the business model was AI services with software as the accelerant.

The pivot started with a phone call from the CEO of a life sciences company who had a problem. His organization had engaged an IT services firm to build AI-powered workflows for its patient safety function. 18 months in, the vendor walked away. The regulatory complexity was too demanding, the accuracy requirements too exacting, and the domain too specialized. When he heard about Graph AI, the CEO reached out.

Raghav took the call and committed: Graph AI would build a working solution. That moment crystallized something the founders had only partially understood. This was a hard domain problem more than it was a hard technical problem. Vijay put it plainly: "We were founders who were technologists at our core, but patient safety is a completely different vertical." So, the team studied how existing systems like Oracle, Veeva, and Aris Global had operated in the territory—and where the boundaries of what software could do ended, and the human labor began.

On the other side of the learning curve, they discovered an opportunity that had been hiding in plain sight. Off-the-shelf LLMs were only delivering roughly 60% accuracy on patient safety tasks, which fell short of the precision a regulated industry could accept. The domain knowledge gap wasn't a blocker; it was the moat. If Graph AI could embed genuine pharmacological expertise directly into the platform, it could deliver near-perfect accuracy. Two and a half weeks later, they delivered an MVP that the previous team couldn't produce in a year and a half. This was how the team's services thesis became a SaaS thesis. And Graph Safety was the result.

## **Building for a domain that doesn't forgive mistakes**

When Graph AI landed its first pharma customer, the team had a working agent platform and a services thesis. At the time, they lacked deep knowledge of pharmacovigilance, so they became students. "One of the challenges we had to overcome very early on was the domain expertise and how to understand the problem and then solve for it," says Vijay.

The founding team learned the workflows from the ground up: how adverse events are reported, how cases are classified, how submissions are structured for the FDA and EMA, and why the existing software had only ever automated a fraction of the process. They hired a medical doctor to join the team and validate the system as it was built, making domain expertise a structural part of their product development process rather than an afterthought. That investment shaped every technical decision that followed.



The most important of those decisions was what kind of AI system to build. The instinct with most AI products is to build a pull model where a user asks a question and the system responds by pulling an answer. Graph AI deliberately went the other direction. "Our system is designed as a push model," explains Raghav. "We push information to the personas that are using our platform—information they need to conduct their workflows. They don't need to trigger an LLM or an agent to get the information they're looking for." In practice, this means the system surfaces the right data, classifications, and recommendations at each stage of the workflow, without requiring users to know how to prompt it. The AI works in the background, and each persona's experience stays familiar.

Underneath that experience is a technical architecture built around two core convictions:

1. Not every problem is an LLM problem. Graph Safety applies deterministic processing where it's appropriate and AI where the problem is genuinely fuzzy—particularly when extracting meaning from unstructured text like clinical trial reports, medical journals, patient records, and regulatory XML feeds.
2. Accuracy in a regulated industry is the only metric that matters. Off-the-shelf LLMs alone, including Gemini and GPT-4, aren't good enough for a process where the wrong classification can result in a missed regulatory submission.

### Closing the accuracy gap

To close the accuracy gap, Graph AI built a knowledge graph that connects adverse events, drug compounds, patient populations, clinical evidence, and regulatory history into a structured, queryable layer that grounds every AI inference in cited source material. Here's how it works:

- When the system surfaces a finding, it shows the user exactly where that finding came from (i.e., which article, which clinical report, which data feed). Hallucination is structurally prevented by requiring every output to be traceable to the source.
- They then fine-tuned their Gemini-based models on top of that knowledge graph by incorporating reinforcement learning from domain expert corrections. When a quality reviewer disagrees with a system inference, that feedback goes back into model tuning.
- The result is accuracy north of 95% in their early deployments.

The validation went further: Google onboarded Graph AI as a Google Cloud partner and is now going to market with Graph Safety to enterprise customers.

Getting the product to production required clearing one more hurdle: regulatory compliance. This is a step most SaaS companies never face. The FDA mandates that any system used in a regulated pharmacovigilance workflow must pass a computer system validation process—a step-by-step manual verification of everything the system does. Graph AI also needed SOC 2 certification and compliance with the FDA's CFR 21 Part 11, which is satisfied by an immutable audit trail and electronic signatures permanently linked to every human-in-the-loop (HITL) sign-off. Beyond these core requirements, the platform was architected to meet a shifting global landscape:



- US FDA AI Guidance: GraphSafety maintains a full Context of Use registry for all 14 AI functions to meet the January 2025 framework.
- EU AI Act: The system satisfies provider obligations for high-risk PV AI through its built-in human oversight and provenance model.
- EMA Reflection Paper: Graph AI provides a formal Vendor Qualification Package and a ready-to-use PSMF template to meet European disclosure requirements.
- Health Canada: A dedicated GCP Canada instance ensures data residency while enabling simultaneous submissions to both the FDA and Health Canada.
- CIOMS Working Group XIV: The architecture aligns with all seven guiding principles, including anchoring outputs to source evidence.

The validation process alone took four months. That rigor became, in retrospect, one of their most durable advantages. Any competitor entering this space faces the same four-month obstacle, but most won't have built their platform from the ground up with auditability in mind.

## **A pricing roadmap that follows the level of automation**

Most enterprise software companies price against other software, but Graph AI prices against the total cost of human labor. The distinction matters because pharmacovigilance has historically been a labor market, not a software one. Traditionally, a pharma company's bill reflects headcount—graduates and reviewers—rather than just licenses. While legacy platforms automate only the software layer, leaving a substantial "labor tax" on top, Graph AI unifies both. By automating the software and the manual work simultaneously, customers achieve a 65% cost savings, delivering a solution that does more for significantly less.

"A lot of these services are delivered by people and less by software," says Raghav. "Today, customers are forced into a model where 75% of their spend goes toward human labor and only 25% toward the software that's supposed to help them. We're inverting that ratio. We've built a platform where 75% of the heavy lifting is handled by software, and only 25% requires human oversight."

Graph AI structures its pricing as a journey that tracks directly how much of the workflow the platform has automated. Today, its two live modules charge per case, or the individual adverse event reports and signals flow into the system before they're classified as confirmed cases. As Graph Safety automates deeper into the process, the pricing unit expands: the next module will charge per periodic report and eventually an enterprise license agreement covering all six planned modules.

This staged model is both commercially pragmatic and philosophically coherent. Each new module automates another layer of human labor, so each new pricing tier reflects a new category of value delivered. By the time a customer is on an enterprise agreement, Graph Safety will have automated roughly 70% of quality review work and 50% of what medical reviewers currently do. The price at that stage will reflect this.



There's also a predictability dimension worth noting. Because pharmacovigilance volumes are relatively forecastable—drug portfolios don't change overnight, and adverse event intake scales predictably with prescribing volume—per-hit and per-case pricing is easier for customers to model than seat-based licensing. It ties Graph AI's revenue directly to what the platform is processing, which aligns incentives cleanly. For founders thinking about pricing in regulated industries, Graph AI's approach offers a useful framework:

**When your product is automating workflows that software didn't touch or didn't do well, don't anchor your pricing in existing software. Anchor it to the value you're unlocking and build your roadmap so each new module earns a new pricing tier by delivering the next layer of capability. The unit of pricing should follow the unit of value.**

## **Graph AI's three routes to market**

Graph AI designed a GTM strategy with three parallel entry points—each targeting a different layer of the pharmacovigilance ecosystem:

1. The first route to market is direct sales to Big Pharma manufacturers, mid-market pharma companies, and CROs. Rather than chasing the largest global enterprises from day one, Graph AI is taking a phased geographic approach—starting with the Americas and Europe, followed by Asia Pacific, and later China. The reasoning is practical: these target markets are currently driving the most significant regulatory shifts in AI-governed safety, requiring a customized product presence.
2. The second motion is more unconventional: turning the existing services industry into a distribution channel. CROs currently deliver pharmacovigilance as a labor-intensive managed service. Rather than positioning Graph Safety as a threat to that business model, Graph AI is offering CROs the platform on which to run their operations—effectively helping them shift to selling software-enabled outcomes. The CROs become resellers with a built-in customer base and a margin expansion story. This is a channel strategy that converts potential competitors into partners.
3. The third is through its Google Cloud partnership. Google has onboarded Graph AI as a GCP partner, acknowledging the precision with which the team has fine-tuned its models for the life sciences domain. Graph Safety will be available on the Google Cloud Marketplace, with Google partnering in the GTM motions, helping bring the product to market. Google's incentive is straightforward: Graph Safety drives meaningful GCP consumption at enterprise scale in a high-value vertical.

Alongside these three motions, Graph AI has been building an advisor network of industry veterans who are actively referring the platform to their pharma and CRO networks. The first cohort of advisors has already generated 15 qualified leads. Graph's GTM is born from a complementary set of skills the founders bring to the table—Raghav with an extensive background in sales, Vijay and Mohan leading the technology and product, and Ashutosh running point on revenue and the company's finances.



## **What drug safety reveals about where AI wins**

Graph AI's story points to something broader about where AI will create substantial value in the enterprise. Vijay describes the platform's long-term potential as a "shift left" opportunity where the same knowledge graph that captures adverse events in the post-marketing phase of a drug's life can feed insights back into earlier stages of drug discovery. What helps companies stay compliant today could tomorrow help them decide which drugs to pursue and which to abandon before spending years and billions finding out the hard way. The platform's compliance infrastructure then becomes a lever for R&D.

But the bigger implication may be for founders building outside of life sciences. The industries where AI will matter most aren't necessarily the ones with the most obvious data or the most sophisticated buyers. It'll be the ones where accuracy requirements are high enough that off-the-shelf models simply fail, where regulatory constraints have kept technology primitive, and where the real incumbent isn't a software vendor but a workforce weighed down by inefficiencies. These conditions produce a problem that's genuinely hard to solve, but genuinely valuable when it is solved.

CASE STUDIES

# Launching AI products that win

Leaders share how they successfully built AI products. [Subscribe to get more insights to help your team win.](#)



Venture insights that matter

[bvp.com/subscribe](https://bvp.com/subscribe)

**About Graph AI:**

Graph AI is building Graph Safety, the industry's first AI-native Patient Safety Operating System, a unified platform purpose-built to re-architect the full adverse event lifecycle and regulatory compliance for pharmaceutical, biotech, and life sciences companies. Founded by Raghavendra Parvataraju (CEO), Vijay Ponukumati (CTO), Mohan Konyala (CPO), and Ashutosh Bordekar (CFO), the company's flagship product, Graph Safety, unifies the full pharmacovigilance workflow across a single, AI-powered platform. Graph Safety is currently used by enterprise customers to monitor the safety profiles of 300 drugs globally. Learn more at <https://graphsafety.ai/>



**Venture insights that matter**

[bvp.com/subscribe](https://bvp.com/subscribe)