IntuitionLabs

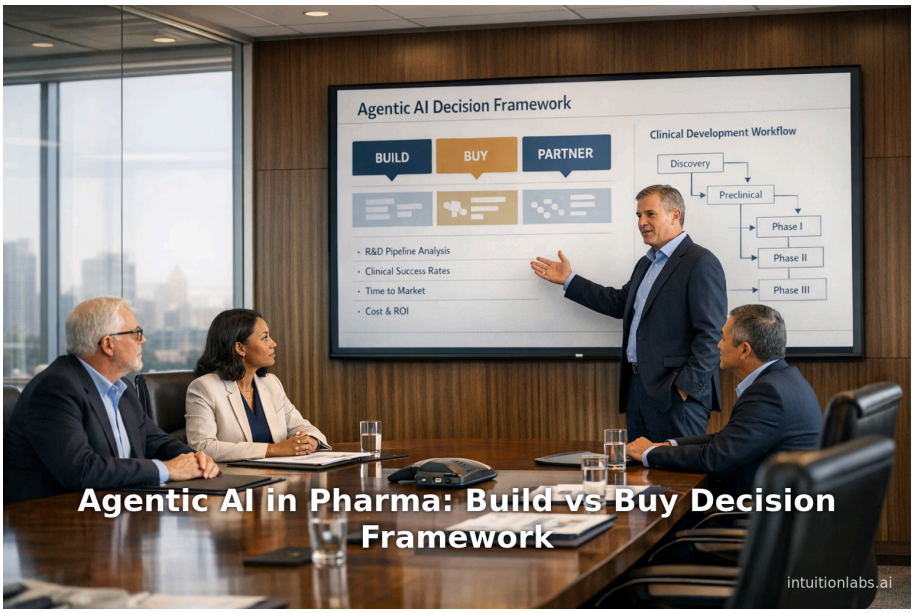
INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Agentic AI in Pharma: Build vs Buy Decision Framework

> Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-04-05 · agentic ai · pharma decision framework · build vs buy ai · pharmaceutical ai · clinical operations ai · drug discovery ai · quality operations ai · life sciences ai

Original article: https://intuitionlabs.ai/articles/agentic-ai-pharma-build-buy-framework



In this article

1. Executive Summary
2. Introduction and Background
3. The Agentic AI Opportunity in Pharma
4. Potential Applications
5. Measurable Impact – Data and Findings
6. Build vs. Buy vs. Partner: Strategic Paths
7. Framework for the Decision
8. In-Depth Analysis: Build, Buy, Partner
9. Building In-House: Innovation vs. Burden
10. Buying Off-the-Shelf: The Vendor Route
11. Partnering: Collaborative Innovation
12. Risks and Considerations
13. Case Studies and Real-World Examples
14. Implications and Future Directions
15. Conclusion

Executive Summary

The pharmaceutical industry stands at a strategic inflection point as **agentic AI**—autonomous, goal-oriented artificial intelligence—begins to move from experimental pilots to enterprise-scale deployment. Agentic AI systems extend beyond conventional AI by planning, executing, and adapting multi-step tasks across complex workflows (often combining large language models with tools, memory, and orchestration (www.fda.gov) (sakaradigital.com)). This shift promises profound impacts across the drug value chain: from accelerating target discovery and clinical trial operations to automating regulated manufacturing and quality processes. Potential benefits in workflows such as deviation investigations should be treated as hypotheses for pilot evaluation. NIST recommends testing AI systems before deployment and regularly during operation (NIST AI RMF).

However, integrating agentic AI also introduces new challenges: **data gaps, regulatory constraints, and talent shortfalls**. Pharma organizations must decide **how** to acquire and deploy these capabilities: to **build** in-house, **buy** off-the-shelf solutions, or **partner** with specialized vendors or consortia. Each path has trade-offs in cost, speed, control, and risk. For instance, building custom AI gives maximum alignment with proprietary data and processes but demands enormous investment (often pushing R&D timelines and budgets), whereas buying established platforms yields quick value but risks vendor lock-in and misalignment with specialized needs (aiireland.ie) (www.causaly.com). A “partner” model—joint ventures or co-development—can bridge these extremes by combining internal assets with external expertise, but requires clear IP/governance agreements (aiireland.ie).

This report provides an in-depth analysis of agentic AI in pharma and a **decision framework** for build vs buy vs partner. We review the **historical context** of AI in life sciences, current state (including market size projections and regulatory developments), and profile AI sourcing strategies across leading firms and consortia. We integrate expert commentary, empirical data, and case studies (e.g. Lilly–Chai Discovery biologics design, GSK–Noetik oncology AI, Salesforce Agentforce deployments at Novartis/AZ) to highlight successes and pitfalls. We also examine a portfolio strategy combining all three approaches, as many experts advocate (aiireland.ie). We conclude with implications for future innovation: how agentic AI may reshape pharma R&D economics, organizational structures, and competitive dynamics in the 2026–2030 horizon.

Introduction and Background

Agentic AI, broadly defined, refers to **AI systems capable of autonomous, goal-directed behavior**. Unlike narrow algorithms or static expert systems, agentic AI can sequence and adapt actions—often via chained planning and tool use—to achieve higher-level objectives without human prompting for each step (sakaradigital.com) (www.fda.gov). This is the next evolution after **generative AI** (which generates content or insights given a prompt). Agentic AI leverages generative models (like LLMs) as components but integrates them with procedure and memory modules to "**act**" on the world: making API calls, retrieving data, manipulating records, etc., all under governance constraints.

In pharmaceutical contexts, agentic AI is poised to transform workflows across research, development, and operations. For example, agentic bots could autonomously triage and synthesize medical literature, design new molecules, orchestrate clinical trial execution, or manage quality investigations end-to-end (sakaradigital.com) (www.drugdiscoverytrends.com). The COVID-era advances in LLMs (e.g. GPT-4, Claude, BioGPT) and tool-augmented agents have set the stage: now companies can envisage AI not just as an assistant, but as a collaborator or even virtual employee performing tasks around the clock.

Historical context: Pharma has used AI (often under terms like "predictive analytics" or "machine learning") for over a decade. Early efforts focused on **predictive modeling** (e.g. identifying biomarkers, optimizing formulations, forecasting demand) (sakaradigital.com). With the 2020s, **generative AI** (LLMs and image generators) became tools for drafting protocols, summarizing literature, or creating synthetic data. Agentic AI marks a third phase. As Sakara Digital (2026) observes, machine-learning in pharma evolved through three generations:

AI Generation	Time Frame	Capabilities	Pharma Examples
Predictive AI	2015–2022	Pattern recognition, statistical modeling on structured data	Predictive maintenance; patient stratification; anomaly detection (sakaradigital.com)
Generative AI	2023–2025	Content generation, summarization, natural language interfaces	SOP drafting; regulatory submission summarization; literature review; chatbots (sakaradigital.com)
Agentic AI	2025–present	Goal-directed multi-step actions; tool integration; adaptive planning	Autonomous deviation investigations; end-to-end batch record review; adaptive clinical trial monitoring (sakaradigital.com)

As shown above, **Agentic AI (2025–present)** transitions AI from answering isolated queries to **pursuing objectives**. A generative model might draft a report; an agentic system would identify data sources, analyze each, generate the report, and route it for review on its own. Critically, agentic systems in pharma must operate within compliance "draft-and-review" loops: e.g. they may compile a deviation root-cause analysis draft which a qualified human reviews (sakaradigital.com). This human-in-the-loop design is both necessity (regulations) and advantage (traceable audit trail).

McKinsey's 2024 generative-AI analysis assigns a modeled **$18–30 billion** annual opportunity to commercial functions in pharmaceutical and medtech companies (McKinsey). Evaluate measured outcomes from agentic AI pilots separately from this projected potential.

These developments occur alongside major infrastructural changes. Foundational AI platform vendors (OpenAI, Anthropic, Google, NVIDIA) have released enterprise-grade agent frameworks; integration firms (MuleSoft, Boomi, Workato) provide connectors to core pharma systems (e.g. SAP, Veeva, LIMS); and regulatory clarity is improving. The FDA, for example, in 2025 issued draft guidance on using AI in product submissions (www.fda.gov) and adopted an **agentic AI tool internally** for review workflows (www.fda.gov). The EU AI Act entered into force on 1 August 2024 and became generally applicable on 2 August 2026, with phased exceptions. The Commission lists 2 December 2027 for high-risk uses in Annex III and 2 August 2028 for high-risk systems embedded in Annex I regulated products. Assess each intended use and product classification rather than treating all pharma AI as high-risk (European Commission).

However, the hype must be tempered with realism. As several analysts warn, AI initiatives often fail due to *capacity and process* mismatches. Gartner (2025) cautions that 40% of agentic AI projects may be cancelled by 2027 due to unclear ROI (www.drugdiscoverytrends.com). Salesforce has emphasized enterprise governance and guardrails for deploying AI agents (www.drugdiscoverytrends.com). Even advanced systems are now understood to require robust orchestration, quality data integration (resolving the "80/20 data problem" of missing context (www.ampcome.com)), and continuous human oversight.

Against this complex backdrop, **pharma executives must craft a decision framework** for agentic AI investment. This framework should consider:

- **Strategic fit:** Is the AI capability core to the company's unique competitive advantage (e.g. novel target discovery leveraging proprietary data)? If so, more internal control (build or close partnership) may be needed. If it's a generic or widely shared need (e.g. sales analytics, standard regulatory drafting), off-the-shelf tools may suffice.
- **Cost and speed:** Estimate project-specific build and operating costs. Causaly notes that recurring costs scale with usage and upgrade cycles (Causaly). Buying can deliver immediate functionality but can entail high subscription costs and less alignment (aiireland.ie).
- **Data readiness:** Pharma data is notoriously siloed (IMS, trial databases, EHRs, R&D ELNs, etc.). Agentic AI thrives on integrated data. Poor infrastructure often dooms build projects, as one consultancy observed (3-month pilot spent 100% time data-wrangling) (www.apparate.com.au).
- **Regulatory risk:** Any agentic deployment must meet traceability, reporting, and audit requirements. This favors approaches where governance is baked in (either via vendor platforms with compliance features or via structured in-house development).

- **Talent and culture:** Specialized “bilingual” experts (biology+AI) are rare (www.pharmexec.com). An external partner can inject skills and best practices, but integrating with corporate culture is critical.

The remainder of this report systematically examines these facets. We first review **current applications of agentic AI in pharma**, illustrating where it delivers value today. We then analyze the **build vs buy vs partner decision**: defining each path, outlining pros/cons, and discussing criteria (with summarized guidelines, see Table 1). We interweave **data-driven insights** and references (e.g. industry surveys, McKinsey/Gartner forecasts) to quantify trends. **Case studies** (real-world partnerships and implementations from 2025–26) highlight lessons, such as Lilly’s collaboration with Chai Discovery for antibody design (creati.ai), Novartis’s adoption of Salesforce’s pharmaceutical AI agent platform (www.drugdiscoverytrends.com), and GSK’s $50M deal with Noetik on cancer digital twins (creati.ai). Finally, we discuss the **organizational and regulatory implications**, and propose future directions (e.g. how a “hybrid” AI strategy can dynamically leverage all three approaches). All major claims are backed by citations to reputable sources spanning industry analyses, technical articles, press releases, and expert commentary.

The Agentic AI Opportunity in Pharma

Potential Applications

Agentic AI can touch virtually every stage of the pharma value chain, but today its **most mature use cases are in manufacturing quality, clinical operations, and commercial analytics** (sakaradigital.com) (www.fiercepharma.com). These areas share common traits: well-defined workflows, large, structured data sets, and high labor intensity—making them prime for automation by an autonomous agent. Key example domains include:

- **Manufacturing / Quality Operations:** As Sakara Digital reports, quality workflows (deviation investigations, batch record reviews, CAPA management) are exploding as first-value targets (sakaradigital.com). For instance, agentic systems can automatically aggregate data from MES, LIMS, CMMS, and QMS to draft deviation root-cause analyses and impact reports. For an investigation pilot, define a baseline and assess cycle time alongside accuracy and review effort. NIST recommends testing AI systems before deployment and regularly in operation (NIST AI RMF). Other tasks like environmental monitoring intelligence and change control are similarly amenable to this approach (sakaradigital.com). In each case, a regulated human review ensures compliance (the agent produces flagged reports rather than final decisions) (sakaradigital.com).
- **Drug Discovery / Preclinical R&D:** The promise of agentic AI here is to compress prohibitively long pipelines. Autonomous “target hunters” could scour literature and patents, propose hypotheses, design molecules via an integrated platform in days rather than months (sakaradigital.com) (creati.ai). Similarly, in lead optimization and toxicology, agents can integrate diverse data streams (bioassays, SAR databases) and suggest experiment designs. For example, in one case study, an agent integrated patient recruitment data and communications to boost trial enrollment by **28%** within a month (www.apparate.com.au). While many discovery projects remain exploratory, pharma–AI partnerships

exemplify the trend: Lilly–Chai Discovery for antibody design (creating purpose-specific antibody models (creati.ai)) and GSK–Noetik for spatial tumor models (creati.ai) suggest that drug design is entering a "deterministic engineering" phase.

- **Clinical Trials:** Agentic AI can automate trial operations (site selection, protocol adherence monitoring, pharmacovigilance alerts) via continuous data integration. Agents monitor incoming EDC (electronic data capture) systems in real time, flag recruitment shortfalls, and even draft safety reports by combining structured data (AE forms) with unstructured clinical notes (sakaradigital.com). This goes beyond retrospective monitoring: agents push notifications to trial teams as soon as issues emerge, akin to an always-on digital CRA.
- **Commercial / Market Research:** In a sponsored article, Customer Insights.AI describes ciATHENA as an agentic and conversational platform with purpose-built market-research analytics applications and natural-language interactions (Customer Insights.AI sponsored article). Treat this as a vendor description: assess task speed, analytical quality and cost against a defined baseline before using the tool in an investment case.
- **Supply Chain & Logistics:** While less mature in life sciences, agentic workflows could optimize global supply networks, managing recalls or shortage events. Agents can coordinate between ERP systems (SAP), quality records, and regulatory constraints to make re-routing decisions or prioritize shipments, acting as "AI supply chain managers." Startups have begun exploring these use cases, although published examples are still limited.

Overall, by automating **high-volume, rules-based human tasks**, agentic AI could liberate skilled scientists and operators to focus on the most critical, creative work. The technology delivers "consistency and availability" beyond human shift limits (artera.io), with the potential to uplift entire processes. As one agentic AI expert puts it, each technical advance that upgrades an agentic workflow (such as adding memory or advanced planning) compounds organizational productivity.

Measurable Impact – Data and Findings

The conjectured benefits of agentic AI are backed by early quantitative indicators. Some key published figures:

- **Adoption Rates:** McKinsey's 2025 cross-industry survey received 1,993 responses from 105 nations, with fieldwork from June 25 to July 29, 2025. Of respondents, 23% said their organizations were scaling an agentic-AI system in at least one business function and another 39% were experimenting. In any individual function, no more than 10% reported scaling agents. These are respondent-level findings, not a pharma-specific pilot rate or the percentage of workstreams in production (McKinsey survey).
- **Efficiency Evaluation:** Assess proposed investigation and record-review gains against a documented baseline, including error rates and human review effort. NIST recommends performance assessment against benchmarks and documented results, rather than assuming that a reported saving transfers to another deployment (NIST AI RMF).

- **Economic Value:** McKinsey's January 9, 2024 report models generative-AI opportunities across the pharmaceutical value chain; its **$18–30 billion** annual range concerns commercial functions, as its May 2024 commercial analysis confirms (January report; commercial analysis). These are modeled opportunities, not realized agentic-AI savings. For context, the global pharmaceutical sales market was ~$1.65 trillion in 2024 (www.fiercepharma.com), so these incremental improvements could meaningfully improve profit margins. In development specifically, Forbes cites Define Ventures data: **80% of large pharma leaders** now prioritize AI to cut therapy development costs (www.forbes.com), and 77% use AI for target ID. Moreover, Define notes "93% of pharma leaders consider medical writing a top AI priority" (www.forbes.com), underscoring that organizational focus heavily leans on near-term ROI use cases.
- **Investment Trends:** Venture funding and partnerships are accelerating. FDA's own investment is a signal (the agency built an Elsa LLM tool for reviewers in 2025 (www.fda.gov), then opened an internal agentic AI challenge in late 2025 to push adoption). CB Insights reports **$3.8 billion** in AI-agent startup funding in 2024 across industries, rather than healthcare alone (CB Insights). The January 2026 discovery announcements illustrate different sourcing models: Lilly's collaboration involves deploying Chai's AI platform for biologics design (Chai announcement), while Boltz explicitly describes a Pfizer partnership to deploy Boltz Lab and its agents (Boltz announcement). Treat the Chai example as a broader AI-sourcing analogy unless an agentic workflow is separately established. Evaluate disclosed deal terms individually rather than inferring an aggregate monetary commitment.
- **Hiring and Workforce:** The pharma sector is visibly scaling its AI talent. In Q2 2024, AI-related job postings in pharma jumped **10%** quarter-over-quarter (www.pharmaceutical-technology.com). Top recruiters in AI roles include Sanofi, AstraZeneca, Novartis, J&J among others (www.pharmaceutical-technology.com). Define Ventures reports that 70% of surveyed pharma leaders viewed AI as an immediate priority, rising to 85% among the top 20 companies (Define Ventures announcement). Yet, specialized talent remains scarce, as noted by PharmExec: hiring specialists (especially "bilingual" R&D domain experts (www.pharmexec.com)) is a critical constraint.

In sum, the data confirm a **high-stakes opportunity** with agentic AI: substantial value is up for grabs **if** deployments are executed properly. Missteps, conversely, risk large sunk costs (a cautionary example described below).

Build vs. Buy vs. Partner: Strategic Paths

Pharmaceutical companies must choose how to obtain agentic AI capabilities. Three broad models dominate:

1. **Build In-House:** Develop AI platforms and agents internally (potentially with external vendors for components) by leveraging corporate data and talent.
2. **Buy Off-the-Shelf:** License commercial AI products or platforms that provide agentic functionality (e.g. SaaS solutions from tech vendors).
3. **Partner/Co-Develop:** Engage in joint development with expert providers, universities, or consortia, sharing resources (and often sharing IP or governance).

Table 1 (below) summarizes key attributes of each approach, drawing on industry analyses (aiireland.ie) (www.causaly.com) (www.drugdiscoverytrends.com).

Approach	Key Advantages	Key Drawbacks	Best Suited When…
Build In-House	• *Control & Customization:* Full oversight of algorithms, data handling, and IP (aiireland.ie) (www.causaly.com). • *Strategic Alignment:* Can tailor to core workflows or proprietary data. (e.g. unique drug targets, internal ELN content.) • *No vendor lock-in:* Ownership of tech and roadmap.	• *High Cost and Time:* Very large upfront R&D investment and long timelines (www.causaly.com) (aiireland.ie). • *Talent Demand:* Requires assembling scarce AI/ML experts plus domain specialists (www.causaly.com) (www.pharmexec.com). • *Maintenance Burden:* Ongoing platform upgrades, governance, and model retraining (hidden TCO) (www.causaly.com).	– The capability is **core** to competitive advantage (e.g. novel modality discovery). – Strict data/control/regulatory needs where vendors can't suffice. – Company can commit long-term capital and sustain an ecosystem.
Buy Off-the-Shelf	• *Speed to Value:* Rapid deployment of mature AI functions; minimal initial setup (aiireland.ie). • *Lower Upfront Cost:* Uses vendor's R&D and infrastructure (OpEx model). • *Proven Solutions:* Benefit from vendor support, updates, and shared risk of development.	• *Limited Customization:* "One-size-fits-all" may not handle niche pharma data or processes fully (aiireland.ie) (www.drugdiscoverytrends.com). • *Vendor Lock-in & Dependency:* Ongoing fees, less control over roadmap, potential IP restrictions. (aiireland.ie) • *Data/Integration:* Integration with legacy systems may still be complex; must fit within agentic pipeline.	– Use cases that are generic or non-differentiating (e.g. AI-driven document summarization, CRM chatbots) (aiireland.ie). – Need early wins to build confidence or when in-house expertise is lacking. – When vendor market is competitive and solution maturity high.
Partner/Co-Dev	• *Best of Both:* Combines proprietary assets (e.g. data, domain knowledge) with partner's expertise and tools (aiireland.ie) (www.causaly.com). • *Shared Risk & Cost:* Costs and technical risk are split; can be faster than pure build. • *Capacity Building:* Opportunity to upskill internal teams through collaboration.	• *Shared IP and Control:* Must negotiate ownership of outputs; may dilute proprietary advantage (aiireland.ie). • *Complex Governance:* Requires clear contracts on roles, timeline, compliance; more moving parts. • *Dependency on Partner:* Could become reliant on partner's roadmap or financial stability.	– Proprietary data or use-case is mission-critical, but internal AI talent is insufficient (aiireland.ie). – Off-the-shelf does not fit due to uniqueness of problem, yet full build is too slow. – Strategic alliances (startups, academia, consortia) are viable to co-develop new AI.

Table 1: Summary of Build vs Buy vs Partner for Agentic AI in Pharma (adapted from discussion (aiireland.ie) (www.causaly.com) (www.drugdiscoverytrends.com)).

As Table 1 shows, there is no one-size-fits-all answer. Leading-edge advice from industry experts is that **the optimal strategy is often a hybrid portfolio approach** (aiireland.ie): "Buy where the use case is generic, build where it is core to [your] competitive edge, and partner where [you] need specialist capability quickly" (aiireland.ie). In practical terms, companies may start with **buying** to prove out lower-risk pilots, then engage partners for more specialized projects, and invest in building only the most strategic agents internally.

Impact on Performance and Risk: This build/buy/partner decision has direct implications for ROI. For example, Causaly emphasizes that in R&D information retrieval systems, one can reach ~60% quality quickly by "spinning up a prototype" with existing tools, but the final ~20–30% improvement (needed for scientific-

grade reliability) is exponentially more expensive (www.causaly.com). This “last mile” strongly favors using proven cores and allocating engineering only to domain-specific edges (the “buy-core, build-edge” approach (www.causaly.com)). Similarly, PharmExec notes that skill and context deficits often make purely in-house builds fall short of “actionable insights” (www.pharmexec.com). On the other hand, purely off-the-shelf projects often hit integration bottlenecks or fail to meet compliance requirements, leading to stalled pilots (www.drugdiscoverytrends.com).

Below we examine these three paths in more depth, drawing on case examples and data.

1. Building In-House Agentic AI

Building agentic AI internally means assembling a cross-functional team (AI/ML engineers, data scientists, integrators, clinical/quality SMEs) and creating a proprietary AI platform tailored to the company’s specific pipelines. The process typically involves:

- **Infrastructure & Data:** Provisioning compute (GPUs/cloud), building a data lake or knowledge graph integrating all internal sources (R&D data, EHRs, lab reports, etc.) (www.causaly.com). Legacy data from past trials, ELNs, and SOPs must be curated and formatted (often a major challenge (www.apparate.com.au)).
- **Core AI Stack:** Selecting or training foundation models (e.g. LLMs or domain-specific models) and building the orchestration layer (agentic AI frameworks) that plans tasks, manages memory, and connects to tools (artera.io) (sakaradigital.com). This “agentic layer” includes components for planning, tool calling, logging, and safety guardrails (www.causaly.com).
- **Evaluation & Governance:** Implementing continuous evaluation systems to ensure accuracy and compliance. For science workflows, this means extensive provenance tracking: systems must link each AI recommendation back to source data (www.causaly.com) (www.fda.gov). As Causaly notes, reproducibility and auditability are non-negotiable (“each answer must trace to sources”) (www.causaly.com).

The **pros** are clear: an in-house agentic system can fully leverage proprietary knowledge (e.g. a pharma company’s unique chemical libraries or patient registries) and bake in domain restrictions from the ground up. Such internal “sector-specific AI” can deliver capabilities unattainable by generic tools. For instance, in 2025 Novartis hired “Agentic AI” directors and integrated AI into its internal data and quality infrastructure (www.novartis.com).

However, the **cons** are steep. As an experienced industry analyst warns, building a sophisticated scientific AI platform is far from just “wiring a connector” (www.causaly.com). The initial prototyping might show 60% solution quality, but pushing scientific applications to the needed 80–90% accuracy (with proofs and citations) is where the real cost and time lie (www.causaly.com). This requires substantial TCO: not just initial development, but ongoing graph maintenance, re-indexing as journals appear, retraining as models evolve, and re-validation whenever a new data source or model is introduced (www.causaly.com). Even factors like taxonomy updates (drug lists, gene names) must be continuously managed (www.causaly.com).

Given these hurdles, build-in-house is most **recommended when**: the AI task is strategic and a competitive differentiator; the data involved is highly proprietary (so even partial vendor solutions are insufficient); and the organization can commit the necessary budget and time horizon. Examples might include deep drug discovery platforms that embed decades of internal data and knowledge. For example, Pfizer's multi-year collaboration with Boltz effectively treated Pfizer's unique preclinical data as an extension of its own build, leading to exclusive high-grade models (creati.ai). Even here, however, Pfizer did not simply "build" from scratch but partnered and leveraged open-source; this hybrid approach underscores how rare is the fully self-sufficient build.

2. Buying Commercial AI Solutions

Buying means **licensing or subscribing to existing agentic AI solutions**. This could be a software platform (cloud or on-premise) built for industry use, or a more packaged SaaS that performs a specific function. Buying drastically shortens time to deployment. One documented commercial sourcing example is Novartis's selection of Salesforce's Agentforce Life Sciences for Customer Engagement (Salesforce announcement).

The **advantages** include: immediate access to mature technology stacks, avoidance of infrastructure and large-scale engineering efforts, and typically an upfront cost lower than building (often an Opex subscription model). If the use case is generic (such as automated call center agents for prescription refills, or extracting info from documents), many vendors exist, making the market competitive and prices reasonable (aiireland.ie). The vendor shoulders much of the R&D and compliance features (e.g. SOC2 certifications, pre-built connectors) which a buyer can leverage.

The **drawbacks** stem from reduced control. Off-the-shelf systems may not natively understand pharma-specific data flows or compliance nuances. For example, Novartis selected Agentforce Life Sciences to support more personalized customer engagement across its global operations (www.drugdiscoverytrends.com). Organizations can configure industry-specific platforms around their governance and connect them with existing systems (www.drugdiscoverytrends.com). Furthermore, reliance on a vendor means ongoing subscription fees and potential lock-in. One must trust the vendor's roadmap and data policies: questions like "does this AI ever train on my data, and how is that secured?" are ever-present.

Thus, buying is best when the target application is **non-proprietary, fast-payoff, or commoditized**. For example, automating standard commercial functions (e-detailing summaries, call log paperwork) or using AI for general safety reporting (with configurable templates). Indeed, the Salesforce Agentforce example illustrates this path: four major pharma (Novartis, AZ, Haleon, Takeda) have all **bought** this API-driven agentic CRM platform to unify their data and automate routine tasks (www.drugdiscoverytrends.com). These companies prioritized getting started quickly with AI agents rather than building a novel solution from scratch.

A practical playbook often advised is "**buy core, build edge**" (www.causaly.com). That is, license a strong base solution for general needs, then invest internal effort (or use partners) to extend it. For instance, a company might subscribe to an agentic analytics tool and then develop custom plugins that connect it to their lab systems or add proprietary knowledge graphs. This hybrid ensures faster initial ROI while still addressing nuances.

3. Partnering and Co-Development

Partnering sits between build and buy. It encompasses several models: contracting a systems integrator or consultancy to build a solution tailored for you; joint ventures with AI-native biotech or software firms; or participating in pre-competitive consortia. The key is **leveraging external expertise and shared resources** without fully ceding control to a third party.

Advantages include gaining specialized skills and reducing time-to-value. A partner likely has frameworks, toolkits, or domain experience that a pharma company lacks. This accelerates development over a pure build. Risk is shared (financially and operationally), and the internal team learns through collaboration. Also, co-development agreements can be crafted so the pharma company retains rights to the developed IP or models.

However, partnering requires **very clear agreements**. Misaligned incentives or vague IP terms can lead to disputes. Governance must prevent “loss of control” risk: for example, who owns improvements made to a jointly built agent? Amid this complexity, pharmaceutical collaborations are common. The Pistoia Alliance’s new agentic AI initiative (launched Sep 2025 with Genentech seed funding) is one illustration of an industry-level partnership to define standards (pistoiaalliance.org), albeit in a pre-competitive sense. For sourcing comparison, the Lilly–Chai and GSK–Noetik arrangements below illustrate broader AI platform and foundation-model sourcing. The Pfizer–Boltz announcement explicitly describes deploying agents. These examples should not all be classified as agentic workflows.

To illustrate, **Table 2 below** lists selected AI sourcing arrangements. The Chai and Noetik entries provide model/platform analogies for the build/buy/partner decision; the Boltz entry documents agent deployment.

Pharma Company	AI Partner	Focus Area	Description
Eli Lilly	Chai Discovery	Biologics Design	Lilly will deploy Chai’s AI platform to design novel biologic therapeutics for multiple targets (Chai announcement). Broader AI platform sourcing example.
GSK	Noetik (AI biotech)	Oncology (NSCLC/CRC)	GSK receives a non-exclusive licence to Noetik’s OCTO-VC foundation models in NSCLC and CRC (Noetik announcement). Foundation-model licensing example.
Pfizer	Boltz (AI lab)	Drug Discovery Agents	Multi-year partnership to put Boltz Lab and its agents in the hands of Pfizer scientists (Boltz announcement). Explicitly documented agent deployment.
Novartis / AZ / Haleon / Takeda	Salesforce (SC Platform)	Commercial Operations	All selected Salesforce’s *Agentforce Life Sciences* platform for enterprise agents. They buy this integrated solution to unify siloed CRM/clinical data and automate tasks (note-taking, trial operations). The platform includes life-sciences-specific capabilities for customer engagement and enterprise workflows (www.drugdiscoverytrends.com).

Table 2: Examples of 2025–2026 AI Sourcing in Pharma (including broader model/platform arrangements and documented agent deployments; sources cited).

These cases illustrate different sourcing models:

- **Platforms, models and agents:** Lilly's use of Chai's biologics-design platform and Pfizer's planned deployment of Boltz Lab agents illustrate distinct technology scopes within the sourcing comparison.
- **Platform licensing:** The Salesforce example is more straightforward "buy", albeit at scale across firms.
- **Model licensing:** Noetik describes its GSK arrangement as licensing AI infrastructure (Noetik announcement).

Crucially, partnership can be a **stepping stone to eventual build**. A company might first partner to understand the agentic domain, then spin up its own team to replicate or extend learned capabilities. Some experts even suggest that "partnering" essentially merges into "build" over time: the collaboration sets up a de facto internal project. This underscores the importance of clear **exit and scaling plans**. As AI Ireland advises, boards should ask: *"Is the partner meant to co-own long-term IP, or will we eventually internalize the tech?"* (aiireland.ie).

Framework for the Decision

Given these nuances, we propose a stepwise decision framework (Figure 1) for pharma executives evaluating agentic AI initiatives:

1. **Define Business Value and Core Needs:** Identify the target use case and its strategic importance. If the task is non-differentiating or easily sourced (e.g. call-center automation), lean toward buying. If it's central to competitive positioning (e.g. proprietary drug design), lean to build or partner.
2. **Assess Data & Talent Readiness:** Auditing data quality, integration, and availability is crucial. The cautionary tale of a mid-size pharma's $3M agentic pilot highlights that *data silos* and legacy formats can derail AI projects (www.apparate.com.au). If core data is scattered or poor, starting with a buy (e.g. for simpler tasks) while investing in data architecture is prudent. Evaluate in-house AI expertise; if the gap is large, partnering or buying with expert support is safer (www.pharmexec.com) (aiireland.ie).
3. **Estimate Costs and Timelines:** Use a Total Cost of Ownership lens. Causaly breaks TCO into building (one-time) vs running (Opex) vs hidden risks (www.causaly.com). Often, quick pilots with purchased tools can secure early wins to build momentum. In parallel, model full build costs and resource needs carefully. If internal R&D budgets and schedules can accommodate a multi-year horizon, a larger build may be justified.
4. **Minimize Risk via Phasing (Portfolio):** Consider starting with buying to get a proof-of-concept and gather user feedback. Then, for higher-impact needs, move to partnerships or in-house development. AI Ireland recommends buying for generic use cases, partnering when specialist capability is needed quickly and building where the capability is core to competitive advantage (aiireland.ie).
5. **Plan for Governance and Change Management:** Regardless of the path, establish strong governance: designate roles (who is "owner" of the AI developments), embed compliance checkpoints, and prepare the organization for change (training stakeholders on AI workflows). The Pistoia Alliance initiative notes that building shared standards and educating stakeholders is key to adoption (pistoiaalliance.org).

By rigorously following such a framework, pharma leadership can turn the technological possibilities of agentic AI into realized value—without succumbing to pitfalls of misalignment or rushed execution. We now turn to deeper explorations of each path, followed by concrete case learnings.

In-Depth Analysis: Build, Buy, Partner

Building In-House: Innovation vs. Burden

When a pharma firm chooses to **build** an agentic AI capability, it commits to developing or rigorously customizing a platform from the ground up. This often involves:

- **Complex Integration:** Multifaceted data ingestion (structured and unstructured) using custom connectors or standards (e.g. FHIR, HL7). For example, in healthcare customer service, experts recommend using model-context protocols (MCP) to let an LLM agent reliably "speak" to EHR or billing systems (artera.io). Similarly, a pharma build project must carefully design its integration layer.
- **Custom Workflow Encoding:** Translating domain workflows and regulations into the agent's reasoning. Causaly emphasizes that generative AI must be integrated into existing workflows rather than treated as a replacement for them (Causaly).
- **Ongoing Evolution:** The AI domain is rapidly changing. Causaly underscores that a build doesn't end at launch—the system will need constant updates: new models, updated ontologies, re-validation, etc. Each upstream tweak (say, a new LLM API) can cascade into many retesting tasks (www.causaly.com).

These demands lead to typical build project failure modes: **scope creep**, ballooning costs, and limited adoption if delivered too late. The *3-million-dollar cautionary tale* illustrates this: a mid-sized pharma sank $3M into agentic drug discovery tools but "the AI models were failing to integrate with existing data systems, leading to a complete standstill" (www.apparate.com.au). The solution required creating a "custom intermediary" to normalize legacy data sources—fixing that alone cut errors 40% (www.apparate.com.au), but only after substantial delay. This anecdote highlights that without early attention to data architecture, build projects can bog down.

Quantifying Build Costs: Use a project-specific total-cost-of-ownership model rather than a universal GPU budget or prototype timetable. Separate infrastructure and model usage, orchestration, integration, specialist staffing, validation, and recurring maintenance. Causaly emphasizes that work continues after the foundation is built (Causaly).

For ROI, model stage-specific cash flows explicitly. Treat any chosen program-spend figure as a scenario assumption, and distinguish avoided expenditure, financing effects, earlier revenue and changes in probability of success. A shorter elapsed development timeline alone should not be counted as an equivalent year of avoided spending.

When Build In-House Makes Sense: Theoretically, build is justified if:

- **Proprietary Advantage:** The AI relies on in-house data that competitors cannot access (e.g. unique biological assays). Then, the solution becomes an asset that others cannot replicate easily.

- **High Regulation or Secrecy:** In cases of extreme data sensitivity or chain-of-custody needs, an internal build may handle compliance more transparently.
- **Long Horizon:** If the organization can afford a long incubation period (e.g. >2 years) and has buy-in from top leadership. This often means the board itself drives it as a “capital strategy” rather than an IT project (aiireland.ie).

Few examples of pure-build agentic projects exist publicly in pharma—most have at least some external involvement. However, firms like Novartis and Roche have internal AI divisions working on proprietary discovery platforms (though results often keep under wraps). The decision between build vs buy turns partly on whether the AI capability is *core to the value proposition*. If the answer is yes, many companies feel compelled to invest internally despite costs.

Buying Off-the-Shelf: The Vendor Route

The **buy** path has accelerated in recent years as specialized AI software providers emerge. Pharma giants are increasingly partnering with technology vendors who offer agentic capabilities as a service. The Novartis–Salesforce Agentforce deal (www.drugdiscoverytrends.com) exemplifies this: Novartis essentially **bought** Salesforce’s life sciences AI platform, a tailored vertical of their agentic CRM tools. The platform promised unification of clinical and commercial data and automated agent workflows (post-call summary drafting, trial site scoping, etc.) (www.drugdiscoverytrends.com). Similarly, companies can purchase analytics hubs where agents synthesize market research data (e.g. the ciATHENA example (www.fiercepharma.com)).

Advantages Revisited: Buying shifts much of the burden to the vendor. The buyer pays licensing or subscription fees (OpEx), often lower than the multi-million CapEx of a build setup. Quick pilots can be launched in weeks, speeding time to initial ROI. Moreover, vendors continuously update their products, so the buyer reaps improvements without direct effort.

Buyer Evaluation: Key considerations include:

- **Vendor Viability:** The vendor must have deep healthcare expertise and stability. A purchase is a bet on the vendor’s long-term roadmap. For example, if choosing a new AI startup for call agents, a pharma firm should vet the startup’s funders and product backlog.
- **Integration and Fit:** Off-the-shelf solutions still require IT integration. Novartis selected Agentforce Life Sciences to support more personalized customer engagement globally (www.drugdiscoverytrends.com). Organizations can then configure deployment around their existing workflows.
- **Cost Structure:** Subscription models mean ongoing costs. Firms must factor net present value of these fees vs build costs. They should also consider vendor lock-in: if a vendor discontinues a feature, the pharma is stuck or forced to adapt.

Confidence and Trust: In some cases, buying also means trusting the vendor’s AI decisions to some extent. This is sensitive in life sciences: How does one verify a vendor’s AI agent is “safe” and “privacy-preserving”? Due diligence often involves requesting evidence of compliance (does the vendor encrypt all data, have

SOC2/ISO27001, etc.) and asking how customer data are used (FDA describes its own internal deployment as using GovCloud models that do not train on input data or regulated-industry submissions (FDA)). Pharma buyers must ensure vendor models are not jeopardizing corporate secrets inadvertently.

Illustrative Example: Beyond Agentforce, consider a hypothetical: a pharma company wants an AI system to monitor social media and literature for adverse event signals. If they **buy** a solution, they might partner with an established pharmacovigilance AI provider. The vendor can quickly ingest social data and run sentiment analysis, flagging potential issues. This avoids the pharma having to build NLP plus databases; however, the pharma must trust that the vendor's model is validated to the regulatory standard. If the system pulls in data incorrectly or misses a safety signal, compliance is at risk. Hence, vendor solutions in regulated AI tend to emphasize transparency: many keep human-in-the-loop checks and provide itemized logs.

Ultimately, buying is in vogue for near-term results. It is often accompanied by the philosophy "start small and prove it first." As AI Ireland advises, **first wins often come from buying** (aiireland.ie). By achieving quick successes in routine use cases, the company can build organizational confidence and refine what features truly need customization before moving on.

Partnering: Collaborative Innovation

Partnerships range from loose collaborations to formal joint ventures. A middle ground, partner engagements can take many forms:

- **Consultant-led builds:** Hiring an integrator or specialist (e.g. Accenture, specialized startups) to build an agentic solution on contract. Here the vendor works for the pharma, but the final IP often belongs to the pharma.
- **Joint development:** Contracts where both parties contribute assets. For instance, Lilly–Chai: Lilly provides data and domain oversight, Chai provides AI platform, and maybe jointly they agree on derivative IP.
- **Model licensing:** Noetik describes its GSK arrangement as licensing AI infrastructure (Noetik announcement).
- **Consortia/Pre-competitive:** Groups of competitors or stakeholders pooling resources. The Pistoia Alliance is pursuing an "Agent-Agent Communication Protocol" to let different AI agents interoperate (pistoiaalliance.org). Although not a direct build, it is a partnership to accelerate the ecosystem.

When Partnering Excels: Generally, partner routes are favored when:

- **Specialized Expertise is Required:** If the needed agentic capability is cutting-edge (say, federated learning for cross-site trials) and no single company has all skills, teaming up is logical.
- **Risk Mitigation:** Shared investment lowers one's exposure. Money and time are split, and there is often external pledging (like commitments to a joint roadmap).
- **Bridge to Build:** Sometimes companies use partners as "transitionary build". They draw on partners to bootstrap an internal program (e.g. initial co-development, then spinning up an in-house team).

Drawbacks and Mitigation: As noted, IP sharing can be contentious. Review data ownership, confidentiality and the handling of future enhancements explicitly in partnership contracts. In *Pharmaceutical Executive*, Rathi Suresh, PhD, describes trust in the vendor and its AI as central to a purchasing decision (Pharmaceutical Executive). Without this, companies risk "vendor-like" dependencies behind closed doors. One tactic is phased agreements: start with pilot contracts (results not IP-transfer), and if successful, move to licensing or separate deals.

We see such strategic partnering in action:

- **Vertical Specialization:** The Lilly–Chai partnership is an example where Lilly recognized it needed Chai's unique "zero-shot antibody" technology (creati.ai). They did not want to build that breakthrough in-house from scratch, but needed a model tailored to Lilly's pipeline. The partnership explicitly includes creating a Lilly-specific model while Chai deploys its platform (creati.ai).
- **Clinical Trial Ops:** Another hypothetical partner pitch could be between a CRO and a tech company. For instance, a company like Medable (which offers endpoints and trial automation tools) could partner with an AI firm to create an agentic platform for patient retention calls and site monitoring, combining Medable's trial data platform with the AI firm's agents. This bundles domain and technical strengths.
- **Pre-Competitive Standards:** The Pistoia Alliance example shows another facet: some partnerships aim to lower overall industry risk. By defining common protocols, multiple pharma firms ensure that in the longer run, the agentic ecosystem is safer and more interoperable. Although not a "paying customer" scenario, it shapes how companies will partner informally (e.g., agreed-upon data standards).

In terms of **output**, partner projects can deliver anything from prototypes to full products. The ultimate goal is generally to transfer capability or product into routine use at the pharma. Contracts should specify these "go-live" criteria and support, to avoid the common situation where a partner tunes an agent but the pharma team never fully integrates it (thereby squandering the investment).

Risks and Considerations

No path is risk-free. Below are critical factors regardless of build/buy/partner choice:

- **Regulatory Compliance:** Agency guidelines are evolving. The FDA's draft guidance (2025) on AI in medical product submissions stresses *model credibility* and 'context of use' definition (www.fda.gov). Any agent used in regulated decision-making (e.g. quality batch release) must maintain evidence of compliance. This implies rigorous documentation, versioning, and human sign-off. Companies must ensure their chosen strategy can satisfy regulatory scrutiny. For instance, an in-house agentic system must be validated like any computerized system under GxP. Vendor solutions should meet GxP standards (some claim to be compliant, but verification is needed).
- **Data Privacy and IP:** Large language models and agent frameworks often raise questions: "did you inadvertently expose confidential data in model training?" or "are reinforcement learning logs archivable?" The FDA specifically stated its agentic models will *not* train on submitted or confidential

industry data (www.fda.gov). Pharma companies need to apply the same caution: their AI vendors should restrict training on sensitive IP or patient data unless explicitly allowed, and should provide isolation (encryption, air-gapped environments) as needed.

- **Talent and Culture:** Building or even adopting agentic AI requires new skills. We repeat that pharma firms typically lack computational-biological bilinguals (www.pharmexec.com). If building, they must recruit aggressively or retrain existing staff. If buying/partnering, they still need "translators" who understand both business needs and the technical concepts. In all cases, cross-functional teams (IT, compliance, R&D, operations) must coordinate; failure here causes governance holes. Organizational change management is often underestimated: real world feedback cycles, usage training, and unrealistic expectations must be managed (as described by Apparate: the initial team expected "instant breakthroughs" but had to reset goals (www.apparate.com.au)).
- **Security:** Agentic systems can act autonomously across systems, raising cybersecurity issues. Attackers could try to corrupt an agent's outputs or hijack its tool actions. Sanitize inputs/commands carefully; implement rigorous identity/role limits. Some vendors (e.g. Salesforce Agentforce) build in "Einstein Trust Layers" for access control (www.drugdiscoverytrends.com), but pharma users often add layers (firewalls, custom parsing) on top. Any strategy must incorporate security—not an afterthought.

Case Studies and Real-World Examples

The following cases illustrate AI sourcing choices, including broader discovery-model arrangements that provide context for the agentic-AI decision framework:

- **Lilly–Chai Discovery (Build/Partner Hybrid):** Eli Lilly's partnership with Chai Discovery (announced Jan 2026) is a hybrid model. Lilly is effectively funding Chai's AI to be specialized on their biologics pipeline. Chai, a deep-tech startup, brings a validated antibody design LLM (Chai-2) with remarkable hit rates (creati.ai). Lilly contributes massive internal biological data sets. The deal is multi-year (focused on biologics targets), aims to yield "purpose-built" models, and implicitly combines vendor licensing with collaborative development. This allowed Lilly to *skip trying to build its own antibody AI from scratch*, and instead co-create with an expert. The result is reported: drug discovery timelines are projected to shrink from **months to mere weeks** for some tasks (creati.ai). This is exactly the accelerated innovation they sought. (No doubt Lilly did due diligence: Chai's Series B funding valuation indicated strong validation of its tech (creati.ai).)
- **GSK–Noetik (Oncology Model Licensing):** In January 2026, GSK obtained a non-exclusive licence to Noetik's OCTO-VC foundation models in NSCLC and CRC. The $50 million covers upfront capital and near-term milestones; annual model-licensing fees are additional (Noetik announcement).
- **Pfizer–Boltz (Open-Source Co-Development):** In early 2026, Pfizer announced a unique model: collaborating with Boltz, an open-source biotech AI lab. Boltz's leadership calls themselves "the Red Hat of biology" (creati.ai): they will refine open models (Boltz-2, BoltzGen) using Pfizer's data. Pfizer retains full IP on any discoveries (creati.ai). This is a partnership that resembles internal build (Pfizer bundles

own historical data) but leverages Boltz's specialized infrastructure. It illustrates that even highly secretive pharma R&D can adopt a semi-open innovation model, gaining from the vibrant open-science community while ensuring results stay proprietary.

- **Novartis and Others with Salesforce Agentforce (Buy):** Novartis committed to a global rollout of Salesforce's Agentforce Life Sciences platform (www.drugdiscoverytrends.com). This vendor solution promised a unified data model across their CRM, making agents a "system of insight" rather than siloed chatbots (www.drugdiscoverytrends.com). Other firms (AstraZeneca, Haleon, Takeda) followed suit in similar deals (www.drugdiscoverytrends.com). This is plainly the "buy" scenario: leveraging a major tech player's agentic tool tailored to pharma. Early use cases include automating documentation and trial operations (e.g. "how to find trial sites?" queries) (www.drugdiscoverytrends.com). Salesforce continues to expand Agentforce Life Sciences with industry-specific capabilities and integrations (www.drugdiscoverytrends.com). Salesforce has emphasized robust guardrails for core processes, including adverse-event detection in calls (www.drugdiscoverytrends.com).
- **FDA Internal AI Deployment:** As an atypical case, the FDA is beginning to use agentic AI in-house. In late 2025, FDA announced that it would roll out pilot agentic AI tools to staff (www.fda.gov). This means compliance roles themselves are benefiting from agents (e.g. automating meeting summaries, inspection scheduling) (www.fda.gov). FDA states that use of the tool is voluntary for its staff (FDA). FDA states that its models operate in a high-security GovCloud environment and do not train on input data or data submitted by regulated industry (FDA). This is an internal deployment safeguard; it does not state that confidential inputs cannot be processed.
- **Pistoia Alliance and Industry Collaborations:** The Pistoia Alliance's agentic AI initiative (launched Sept 2025) is a broader, collaborative example (pistoiaalliance.org). Major pharma (Genentech, others) are funding a pre-competitive effort to define **standards and frameworks**, such as an "agent-agent communication protocol" (pistoiaalliance.org). While not directly about build/buy/partner choices for a product, it underscores that even competitors recognize the need to partner on foundational infrastructure. Participating in such consortia can be considered a long-horizon "partner" strategy that may not pay immediate returns but shapes the landscape (e.g. if the protocol becomes an industry standard, firms save future integration costs).

Across these stories, common **lessons** emerge:

- **Integration Matters:** Every example shows that an agentic AI without seamless data integration or aligned governance falls flat. Successful deployments invest heavily in middleware and oversight.
- **Human Oversight Remains Essential:** All agents operate with a human in the loop (especially in high-stakes tasks). This is regulatory necessity and good practice. Thus, every model needs a QA phase and must log its reasoning (strengthening adoption).
- **Domain Specificity is Key:** Generic LLMs alone are seldom enough. High-performing systems refine models with domain data and sometimes custom AI architectures (as seen with Chai-2, OCTO-VC, BoltzGen). This often dictates partnering with domain-tech specialists.

- **Portfolio Strategy in Action:** None of these cases is purely build or pure buy. The Agentforce example illustrates commercial platform adoption within a broader portfolio strategy. The partnerships are building at the core plus buying platform components. Companies are already applying the “buy, then partner, eventually build” philosophy (aiireland.ie).

Implications and Future Directions

Organizational and Cultural Change

As agentic AI matures, organizations must adapt beyond technology. Key imperatives include:

- **Governance & Oversight:** Establish AI governance committees involving legal, compliance, R&D, and business leaders. Frame agentic AI as an enterprise process, not an isolated IT project. For instance, boards should add AI sourcing decisions to the risk register (as recommended in strategic guidelines (aiireland.ie)).
- **Skill Development:** Build internal expertise in “AI stewardship.” This includes data engineers, model validators, and LLM prompt-calibrators. Encourage reskilling managers in data literacy (the so-called “provenance literacy”).
- **Regulatory Liaison:** Engage early with the FDA/EMA on agentic use cases. Submit AI-related proposals in early interactions. In its announcement, FDA encourages sponsors to consult on AI model credibility (www.fda.gov). Proactively aligning with emerging guidance (e.g. on AI transparency (www.europeanpharmaceuticalreview.com)) can avert future compliance headaches.
- **Risk Management:** Treat AI procurement as a fiduciary decision (aiireland.ie). That means rigorous ROI projections, portfolio analysis, and pilots measured by key performance indicators (speed, error rates, user satisfaction). Transparent reporting on AI outcomes avoids the “no measurable impact” trap some CEOs see (aiireland.ie).

Market and Ecosystem Evolution

Looking ahead, several trends will shape how agentic AI decisions play out:

- **Maturing Standards:** As proposed by Pistoia, we may soon see industry standards for AI agentry (communication protocols, benchmarks). These will lower integration barriers and reduce vendor lock-in if widely adopted.
- **AI-Ready Data Platforms:** We anticipate a wave of investments into robust data fabrics and knowledge graphs. If a firm has ready “context graphs,” it can plug agent frameworks in much faster (www.causaly.com). Those who fail to modernize data will see their agent projects stall early.
- **Hybrid Architectures:** Following AI Ireland’s view, the future is “build AND buy and partner” (aiireland.ie). For pharma, this might look like an internal “AI platform shell” that uses both open-source and vendor modules, integrated via cloud-native microservices.
- **Increasing Competition:** The value of agentic AI may also heighten competition for talent and startups. VCs are already pouring money into domain-specific AI biotechs (Biomed Nexus lists dozens of new AI-driven drug companies in 2026). Smart pharma firms might hedge by investing in or acquiring promising AI players, thereby converting a partnership into an acquisition if alignment is strong.

- **Policy and Ethics:** Finally, public trust and ethics will shape agentic AI. Will patients and healthcare authorities accept decisions that are “agent-assisted”? Governance must address biases (the EU AI Act emphasizes fairness) and ensure patients’ rights (especially with personal data). Pharma companies seen as pioneers must balance innovation with transparency—an ethical stance that double as a long-term competitive advantage.

Conclusion

Agentic AI holds transformative potential for the pharmaceutical industry, promising to automate complex tasks and accelerate innovation. However, capturing that potential requires strategic choices. The **Build vs Buy vs Partner framework** provides a lens for pharma leaders to allocate resources effectively, balancing speed, cost, and control. Our analysis shows that *no single path suffices for all use cases*. The recommended strategy is often a hybrid portfolio: use **commercial tools** to get quick wins in generic areas; engage in **targeted partnerships** for specialized pipelines; and selectively **build in-house** for capabilities that truly define one’s competitive edge (aiireland.ie).

Throughout, the overriding theme is **alignment**: alignment of AI initiatives with core business objectives, alignment of algorithms with domain expertise, and alignment of technical execution with regulatory and ethical requirements. These alignments should be tested against deployment-specific outcomes. NIST recommends testing AI systems before deployment and regularly during operation (NIST AI RMF). Conversely, misaligned projects risk being discarded and even eroding confidence in AI.

The pharmaceutical ecosystem itself is adapting. Collaborative projects (like the Pistoia Alliance agentic standards) and cross-industry partnerships (e.g. biotech–AI tie-ups) are laying the groundwork for an agentic future that can fulfill the lofty aim of bringing safer, more effective medicines to patients more efficiently. By carefully evaluating build/buy/partner decisions in light of evidence and real-world examples, pharma can navigate this evolution and build the AI-driven capabilities that will define leadership in the coming decade.

References: The claims and data above are drawn from industry reports, press releases, and expert analyses (e.g. Fierce Pharma, FDA announcements, McKinsey and Gartner studies) as cited throughout. All statements of fact are accompanied by source notes (aiireland.ie) (sakaradigital.com) (www.forbes.com), with source URLs included in the inline links.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://www.fda.gov/news-events/press-announcements/fda-expands-artificial-intelligence-capabilities-agentic-ai-deployment
2. https://sakaradigital.com/blog/agentic-ai-life-sciences-pharma-biotech-pilots-production-2026
3. https://airc.nist.gov/airmf-resources/airmf/5-sec-core
4. https://aiireland.ie/2026/03/20/build-vs-buy-vs-partner-the-ai-decision-every-c-suite-must-get-right-in-2026
5. https://www.causaly.com/blog/building-vs-buying-agentic-ai-what-executives-need-to-know
6. https://www.drugdiscoverytrends.com/novartis-picks-salesforces-agentforce-as-pharma-bets-big-on-ai-agents
7. https://www.mckinsey.com/industries/life-sciences/our-insights/early-adoption-of-generative-ai-in-commercial-life-sciences

8. https://www.fda.gov/news-events/press-announcements/fda-proposes-framework-advance-credibility-ai-models-used-drug-and-biological-product-submissions
9. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
10. https://www.ampcome.com/post/agentic-ai-healthcare-companies
11. https://www.apparate.com.au/blog/agentic-ai-pharma
12. https://www.pharmexec.com/view/to-build-or-to-buy-determining-your-ai-path
13. https://creati.ai/ai-news/2026-01-23/pharma-ai-platform-investments-partnerships-drug-discovery
14. https://www.fiercepharma.com/sponsored/agentic-ai-and-future-pharma-market-research
15. https://artera.io/blog/agentic-ai-for-heathcare
16. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2025
17. https://www.mckinsey.com/industries/life-sciences/our-insights/generative-ai-in-the-pharmaceutical-industry-moving-from-hype-to-reality
18. https://www.forbes.com/sites/sindhyavalloppillil/2025/07/17/ai-in-pharma-era-where-big-tech-leads-startups-scale-and-incumbents-strategize
19. https://www.cbinsights.com/research/ai-agent-market-map
20. https://www.businesswire.com/news/home/20260108131261/en/Chai-Discovery-Announces-Collaboration-with-Eli-Lilly-and-Company-to-Accelerate-Biologics-Discovery
21. https://www.boltz.com/launch
22. https://www.pharmaceutical-technology.com/jobs/artificialintelligence-hiring-pharmaceutical-industry
23. https://www.prnewswire.com/news-releases/big-pharmas-ai-adoption-is-reaching-a-tipping-point-amid-mounting-pressure-and-strategic-opportunity-finds-new-report-from-define-ventures-302507676.html
24. https://www.novartis.com/be-nl/careers/career-search/job/details/req-10072907-associate-director-agentic-ai
25. https://investor.salesforce.com/news/news-details/2025/Salesforces-Agentforce-Life-Sciences-Selected-by-Novartis-to-Drive-More-Personalized-Customer-Engagement-Globally/default.aspx
26. https://pistoiaalliance.org/ai/pistoia-alliance-unveils-agentic-ai-initiative-and-seeks-industry-funding-to-drive-safe-adoption
27. https://www.businesswire.com/news/home/20260108468293/en/GSK-Licenses-Noetiks-AI-Foundation-Models-in-Anchor-Partnership-to-Transform-Cancer-Therapeutic-Research-and-Development
28. https://www.europeanpharmaceuticalreview.com/ai-act-data-governance-and-compliance-strategy-implications-in-pharma/264445.article

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore AI enablement, pharma and biotech software development, data engineering and BI, and Veeva services. Contact IntuitionLabs to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: https://intuitionlabs.ai

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.