

EBOOK



The Governed Agent.

A CxO Playbook for Scaling Agentic AI Across Pharma, Biopharma and Medtech

Built on AWS, the Databricks Lakehouse and Anthropic's Claude, grounded in Veeva Vault data.

kasmoglobal.com

The Agentic AI Moment: From Experiments to Enterprise Risk

Every commercial, clinical, R&D and manufacturing function in pharma, biopharma and medtech is already experimenting with AI copilots and point tools. The question for 2026 isn't whether to adopt agentic AI — it's whether your organization can trust what it does once it's live.

Nearly eight in ten life sciences companies use generative AI today — yet roughly 80% of them report no tangible bottom-line benefit from it.¹ Agentic AI, systems that reason, plan and act rather than simply respond, is where that promise is supposed to change. It's also where the risk concentrates: handing an autonomous system real decisions, without the governance to back it up, is how a promising pilot becomes an expensive, cancelled program.

40%+

Gartner predicts **more than 40% of agentic AI projects will be canceled by the end of 2027**, due to escalating costs, unclear business value and inadequate risk controls.²

45%

of medtech executives say agentic AI is already in production somewhere in their organization.

7%

of those deployments are described as fully embedded and audit-ready.³

The gap between those two numbers is the whole story: adoption has outrun governance. Leaders aren't short on pilots — they're short on agents they can defend to a regulator, an auditor, or a patient.

In This Industry, 90% Accurate Isn't a Success

In most industries, an agent that gets it right nine times out of ten is a productivity win. In pharma, biopharma and medtech, a 90%-accurate agent making a clinical, regulatory or supply chain call is a liability, not a milestone.

Every agent your organization deploys eventually has to answer a hard question: which clinical protocol, which validated dataset, which manufacturing event led to this output? If the answer is “we’re not sure,” that agent doesn’t belong in production — no matter how capable the underlying model is. That is the real reason most agentic AI pilots in this industry stall before scale: not a shortage of intelligence, but a shortage of lineage, governance and domain grounding.

“The industry doesn’t have an AI-capability problem. It has a trust problem — and trust is an architecture decision, not a model decision.”

The upside is real for leaders who solve the trust problem. McKinsey’s analysis of 270 workflows across pharmaceutical and medtech companies puts a number on what’s at stake over a 3–5 year horizon:⁴

PHARMA & BIOPHARMA

- 5–13 pt growth acceleration
- 3.4–5.4 pt EBITDA margin improvement
- 75–85% of workflows contain agent-ready tasks

MEDTECH

- 3–7 pt growth acceleration
- 2.2–4.7 pt EBITDA margin improvement
- 70–80% of workflows contain agent-ready tasks

Figures reflect industry-wide modeling, not a specific Kasma Global client result.

⁴ McKinsey & Company, *Reimagining life science enterprises with agentic AI* (September 2025). Based on 270 workflows, 1,200+ tasks and 180 job families.

The KasmO Approach: A Governed Digital Workforce

We build agentic AI the way life sciences and medtech organizations actually need it built — grounded in your own enterprise data, governed end-to-end, and designed alongside your domain experts, not retrofitted from a horizontal chatbot framework.



Governed data foundation

AWS + the Databricks Lakehouse (Delta Lake, Unity Catalog, MLflow) — every agent decision carries lineage back to its source.



Domain-grounded reasoning

Anthropic's Claude, including Claude for Life Sciences connectors into the scientific and clinical tools your teams already use.



One source of truth

Veeva Vault-anchored data alongside sources like IQVIA, Komodo Health and Cardinal Health — unified instead of stitched together.



Human-in-the-loop by design

Every KasmO agent is advisory. Qualified humans keep final sign-off on clinical, regulatory and commercial decisions.

On that foundation, we've built four solutions that turn a specific life sciences or medtech function into a governed digital workforce:

01 · COMMERCIAL

Commercial
Copilot

02 · CLINICAL

Clinical Trial
Intelligence Agent

03 · R&D

R&D
Agents

04 · SUPPLY CHAIN

Supply Chain &
Manufacturing
Intelligence

01 Commercial Intelligence & Engagement

COMMERCIAL COPILOT

Commercial teams sit on some of the richest, most fragmented data in the enterprise — Veeva CRM, claims, patient support, and distributor data that rarely talk to each other. Commercial Copilot unifies that data into a single 360° view of patients, providers and accounts, then layers agentic reasoning on top: natural-language Q&A over governed KPIs, next-best-action recommendations, and always-on monitoring of engagement signals.

USE CASES

- Natural-language Q&A over 30+ governed commercial KPIs
- Next-best-action recommendations for HCP engagement and rep prioritization
- MLR-grounded call-plan and content-personalization support
- Real-time adherence and engagement signal monitoring

BUSINESS OUTCOMES

- ▲ **4–8%** commercial revenue increase and **5–9%** spend reduction from agentic commercial workflows.⁵
- ▲ **15–25%** reduction in rep administrative burden, freeing selling time.⁵
- ▲ **20–40%** HCP engagement lift reported by comparable next-best-action deployments in life sciences.⁶

CLIENT SNAPSHOT · SPECIALTY PHARMA COMMERCIAL TEAM

Patient and provider data was scattered across Veeva Vault CRM, claims, and distribution systems — so by the time a report was ready, the best engagement window had passed. Commercial Copilot unified the data into one governed 360° view, with AI/BI Genie letting the team ask plain-language questions and MLflow scoring the accounts most likely to convert.

“Life sciences commercial teams don’t need another dashboard — they need one trusted, governed source of truth that can also tell them what to do next.” — Kasmo Global

02

Clinical Development Acceleration

CLINICAL TRIAL INTELLIGENCE AGENT

A single protocol amendment can cost months and millions; a single mis-ranked trial site can cost a program its timeline. The Clinical Trial Intelligence Agent is a supervisor agent plus six specialists — study design, site selection, investigator evaluation, enrollment & diversity, competitive landscape, and evidence & reporting — producing ranked, evidence-cited recommendations your team can defend to a steering committee. Every recommendation is advisory-only, with mandatory human clinical review before anything ships.

USE CASES

- Evidence-cited trial site and investigator ranking and forecasting
- Protocol design support and amendment risk flagging
- Enrollment and diversity/access forecasting
- Competitive trial landscape monitoring

BUSINESS OUTCOMES

- ▲ Up to **50% faster** trial design with **25% fewer** protocol amendments.⁷
- ▲ **\$400M+ NPV gain** associated with a 12-month reduction in trial timelines for a typical late-stage program.⁷
- ▲ **45–50%** time savings in biostatistics and data-management workflows.⁷

CLIENT SNAPSHOT · MID-SIZE CLINICAL-STAGE BIOPHARMA SPONSOR

Site and investigator selection meant manually stitching together CTMS, EDC and public sources like ClinicalTrials.gov and NIH RePORTER — a multi-week exercise leaning on institutional memory. The Clinical Trial Intelligence Agent unified internal and public data into one governed lakehouse and turned it into evidence-linked, explainable rankings in minutes.

“Clinical operations teams don’t need an AI that makes decisions for them — they need one that does the research faster and shows its work.” — Kasmo Global

03 R&D & Scientific Discovery

R&D AGENTS

Your scientists, not your engineers, are your most valuable AI asset — but most of their time still goes to literature review, protocol drafting and reconciling data across a dozen systems by hand. R&D Agents connect directly into the tools your research teams already use — Benchling, PubMed, Scholar Gateway/Wiley, BioRender, Synapse.org and 10x Genomics — through Anthropic’s Claude for Life Sciences connectors, live in weeks rather than months of custom integration work. Every output carries a mandatory scientist-review workflow before it informs a decision.

USE CASES

- Automated literature synthesis and adverse-event summarization
- Protocol and regulatory-document drafting support
- Bioinformatics QC and multi-omic data reconciliation
- Clinical document and patient-note classification

BUSINESS OUTCOMES

- ▲ **20–30%** reduction in early drug discovery timelines.⁸
- ▲ **21–30%** of R&D capacity freed for higher-value scientific work.⁹
- ▲ Regulatory submission prep compressed from **weeks to hours** in leading deployments.⁹

CLIENT SNAPSHOT · CLINICAL-STAGE BIOTECH RESEARCH ORGANIZATION

Years of experiments and protocols lived in Benchling — rich, well-organized, and out of reach for any general-purpose AI assistant. Using Claude for Life Sciences’ pre-built connectors instead of custom integration work, Kasmo Global stood up literature-review, protocol-drafting and bioinformatics-QC agents grounded in the client’s own lab data — in weeks, not months.

“The hard part of R&D AI was never the model — it was months of integration work just to get an agent to see a scientist’s own data.” — Kasmo Global

04

Manufacturing, Supply Chain & Quality

SUPPLY CHAIN & MANUFACTURING INTELLIGENCE PLATFORM

For medtech and biopharma manufacturers alike, supply chain and quality are where agentic AI is already furthest along — Deloitte’s 2026 medtech survey found manufacturing and supply chain among the very first functions to reach production.³ But “in production” doesn’t mean “audit-ready”: the same survey found only 7% of medtech agentic AI deployments are fully embedded with proper validation. Kasmo’s platform is built on Lakebase — a unified OLTP/OLAP architecture that streams ERP, supplier, inventory and distribution data in real time, with no separate ETL pipeline — and keeps Unity Catalog-governed lineage and CQA tracking for GxP-regulated manufacturing, including biologics and cell/gene therapy.

USE CASES

- Real-time OTIF monitoring and autonomous logistics re-routing
- ML-based demand forecasting and dynamic safety-stock optimization
- Manufacturing anomaly detection and equipment-use optimization
- Digital twins and predictive maintenance for production lines and bioreactors

BUSINESS OUTCOMES

- ▲ **20–30%** lower inventory levels through ML-based demand forecasting.¹⁰
- ▲ **50–55%** time savings in supply chain planning workflows.⁹
- ▲ **15–20%** of medtech operations capacity freed for higher-value work.⁹

CLIENT SNAPSHOT · GLOBAL BIOPHARMA MANUFACTURER (BIOLOGICS & CELL/GENE THERAPY)

ERP, supplier, inventory and distribution data each lived in its own system, so shortages typically surfaced only after they’d already hit a shipment. Streaming that data directly into Lakebase — no separate ETL pipeline — put real-time OTIF tracking, shortage prediction and predictive maintenance in front of the team for the first time.

“Supply chain and manufacturing teams don’t need another dashboard that tells them what already went wrong — they need to see it coming.” — Kasmo Global

The Governance Backbone

Deploying an agent is easy. Deploying one you can defend to a regulator, an auditor or a patient is not. Every Kasmo agent is built against four non-negotiables.

01**Lineage by default**

Unity Catalog traces every agent output back to the specific dataset, document or manufacturing event it referenced.

02**Human-in-the-loop where it matters**

Agents are advisory by design. Final clinical, regulatory and commercial decisions stay with a qualified human.

03**Any model, one governance layer**

Your choice of reasoning engine — governed by the same access controls, audit trails and monitoring.

04**Built by your domain experts**

Low-code agent-building lets scientists, clinicians and analysts refine agents directly — instead of a six-month engineering backlog.

This is the gap Deloitte's medtech survey put a number on: 45% of leaders have agentic AI in production, but just 7% call it fully embedded and audit-ready.³ Closing that gap — not shipping another pilot — is the actual competitive advantage available in 2026.

³ Deloitte Center for Health Solutions, medtech leader survey (2026).

Why Life Sciences and Medtech Leaders Choose Kasmoo

Generic AI platforms retrofit a horizontal framework onto pharma, biopharma and medtech. Kasmoo Global was built the other way around.

Pharma-exclusive, not a vertical add-on

Every solution is designed specifically for commercial, clinical, R&D or manufacturing functions in life sciences and medtech — not a general enterprise agent framework relabeled for the industry.

Named products, not bespoke projects

Commercial Copilot, the Clinical Trial Intelligence Agent, R&D Agents and the Supply Chain & Manufacturing Intelligence Platform are shipped solutions you can evaluate today, not a multi-month scoping exercise.

Weeks to value, not months of integration

Pre-built connectors into Benchling, Veeva Vault, PubMed and the rest of your existing stack mean agents go live in weeks.

Governance as the differentiator, not the disclaimer

Explainable scoring, mandatory human review and end-to-end lineage are built into the architecture — the exact gap most agentic AI vendors gloss over.

Deploys in your environment

Solutions run inside your own AWS/Databricks workspace and cloud of choice — your data never has to leave your governance boundary.

Turning Agentic AI Into Measurable Impact

Pharma, biopharma and medtech organizations are entering a decade where intelligent systems don't just assist people — they actively collaborate in discovering new therapies, running trials, engaging providers and keeping supply chains resilient.

The hype cycle is peaking, but the leaders are already looking past it. They aren't experimenting with chat; they're deploying agents that act — and that they can prove are acting correctly. The difference between the 40%+ of agentic AI projects that will be canceled and the leaders who scale won't be the model they chose. It will be whether governance was built in from day one.²

With Kasmoo Global, you're not deploying a model — you're onboarding a governed digital workforce that understands your science, respects your compliance requirements, and gets sharper with every review cycle.

Start Your Agentic AI Readiness Assessment

See which of your commercial, clinical, R&D or manufacturing workflows are ready for a governed agent — and which need a data foundation first.

TALK TO KASMO GLOBAL

