Part 05 STACKAI WHITEPAPER

AI Agents for Medical Device and Life Sciences

How medical-device makers, diagnostics companies, CROs, and nutraceutical brands are deploying AI agents across order operations, regulated complaints, claims-compliant product guidance, and research analytics, without compromising FDA defensibility or IP.

A practical guide for VPs of Operations and Quality, Regulatory and Compliance leaders, Commercial Operations, and Heads of Clinical and Research Analytics.

PREPARED BY

Shani Fargun

VP Healthcare at StackAI

sfargun@stack-ai.com

Executive summary

Life sciences and medical technology sit under the strictest regulatory scrutiny in the economy, and that scrutiny lands squarely on the operational and commercial work: the FDA-regulated complaints process, the structure-versus-function line on product claims, the validation and traceability of any system that touches quality, and the confidentiality of clinical, contractual, and IP data. Generic AI tooling ignores all of it. Regulated companies need agents that are governed, well-sourced, model-agnostic, and defensible.

This whitepaper covers five use cases proven in production across a life-sciences instrumentation company, a clinical research organization, and a compliant supplement manufacturer: order operations from case to shipment, FDA-regulated complaints intake, claims-compliant product and protocol guidance, research and contract analytics, and commercial account intelligence. Each is built to satisfy the regulatory posture that defines the industry: an auditable trail for quality-relevant actions, claim guardrails that keep marketing and clinical content on the right side of the line, and single-tenant deployment so proprietary data never leaves your control.

The companies that operationalize this will move faster on the repetitive, documentation-heavy work that consumes their teams, while strengthening rather than weakening their regulatory position.

02 THE PROBLEM

The regulated-operations problem

The work that consumes life-sciences and medtech teams is repetitive, documentation-heavy, and regulated at every step.

- Order and shipment operations require rekeying cases into the ERP, generating advance shipment notifications, and handling the extra customs, HTS, and country-of-origin requirements of international shipments, all with a transcription-error rate that drives corrections downstream.
- Complaints handling is FDA-regulated. Every inbound signal has to be reviewed, classified, and captured as a structured complaint record that drives regulatory reporting, and it has to be defensible under audit.
- Product and protocol guidance for supplements and nutraceuticals lives under DSHEA, where structure and function claims are permitted but disease and treatment claims are not, and a single non-compliant statement is a regulatory problem.
- Research, contract, and analytics work is buried in MSAs, study contracts, timesheet and utilization data, and multi-million-dollar syndicated datasets that only a small analyst team can currently query.

None of this is a good fit for ungoverned AI. All of it is a good fit for governed agents that draw on approved sources, that keep a human on anything consequential, that log everything, and that run where your IP stays yours.

03 PLATFORM FIT

Why life sciences needs different AI infrastructure

CHARACTERISTIC	WHY IT MATTERS FOR PLATFORM SELECTION
FDA and quality defensibility	Any agent touching complaints or quality-relevant records must produce a traceable, auditable record that stands up in an FDA inspection.
Claim guardrails	For supplements and nutraceuticals, the platform must enforce the DSHEA structure-versus-function boundary and route disease or treatment questions to a human reviewer.
IP and data confidentiality	Clinical data, executed contracts, CVs, and proprietary datasets cannot leak. Single-tenant deployment and no model training on your data are prerequisites.
Model-agnostic and validatable	Regulated buyers need to choose and change models, run head-to-head evaluations, and document why an output is reliable, not be locked to one vendor's model.
Enterprise system reach	The work lives in Salesforce, Rootstock and other ERPs, SharePoint, Snowflake, Workday, and QMS platforms. The platform has to orchestrate across all of them.

04 THE FIVE USE CASES

The five use cases

Five use cases are proven in production across a life-sciences instrumentation company, a clinical research organization, and a compliant supplement manufacturer.

1

Order Operations: Case to Sales Order to Shipment

Converts inbound Salesforce cases into ERP sales orders and generates domestic and international advance shipment notifications.

2

FDA-Regulated Complaints Intake

Classifies inbound signals as complaints, extracts the regulated fields, and creates the structured complaint record that drives reporting.

3

Claims-Compliant Product and Protocol Guidance

Answers product and protocol questions inside DSHEA guardrails and builds personalized, compliant protocols.

4

Research, Contract, and Utilization Analytics

Natural-language access to project, timesheet, contract, and syndicated research data.

5

Commercial Account Intelligence

Generates business and quarterly reviews and drafts research-backed RFP and proposal responses.

1 Order Operations: Case to Sales Order to Shipment

THE PROBLEM

In an instrumentation or device business, an analyst reads each inbound Salesforce case and rekeys it into the ERP as a sales order, then generates the advance shipment notification, in the right format, with the extra documentation international shipments demand. It takes hours, and every rekey is a chance for a transcription error that turns into a shipment correction.

HOW AN AI AGENT HANDLES IT

An agent classifies each inbound case (sales order, complaint, RMA, PO request, spam) and routes it. For sales orders it extracts the customer, SKUs, quantities, ship-to, and PO reference, validates against pricing and eligibility rules, and creates the sales-order record in the ERP (for example Rootstock), writing the case number and any exceptions back to Salesforce for visibility. Companion flows generate advance shipment notifications for domestic and international shipments in the format each customer or carrier expects (EDI 856, PDF, custom template), with the international variant handling customs documents, HTS codes, country-of-origin, and hazmat declarations. Routine cases flow through without a human touch; low-confidence cases route to a queue.

WHY IT MATTERS

It compresses the case-to-order cycle from hours to seconds and eliminates the transcription errors that drove shipment corrections, while keeping a human on the exceptions.

BUILT FOR	Customer Success, Order Operations, Supply Chain
KEY KPIs	Case-to-order cycle time, transcription-error rate, shipment-correction rate
TYPICAL IMPACT	Hours to seconds on routine cases, near-zero transcription error, humans focused on exceptions

2 FDA-Regulated Complaints Intake

THE PROBLEM

Medical-device and life-sciences companies must run an FDA-regulated complaints management process. Every inbound email and case has to be reviewed, classified, and captured as a structured complaint record, which then drives regulatory reporting. Manual intake is slow and produces inconsistent records that are hard to defend in an audit.

HOW AN AI AGENT HANDLES IT

An agent reviews inbound emails and Salesforce cases, classifies each as a complaint, extracts the product, lot number, complaint description, severity, and adverse-event indicators, and creates the structured complaint record in the QMS (for example ComplianceQuest), which then drives downstream regulatory reporting. A human reviews and approves before the record is finalized.

WHY IT MATTERS

It dramatically reduces the QA team's manual intake time and, more importantly, improves the consistency and defensibility of complaint records for FDA audit, which is where inconsistent manual intake creates the most risk.

BUILT FOR	Quality Assurance, Regulatory Affairs, Complaints Management
KEY KPIs	Intake time per complaint, record consistency, audit-readiness
TYPICAL IMPACT	Faster intake, consistent and defensible complaint records, reduced audit risk

3 Claims-Compliant Product and Protocol Guidance

THE PROBLEM

For supplement and nutraceutical companies, reps, customer success teams, and licensed practitioners need fast, accurate product and protocol answers, but every answer has to stay on the right side of the DSHEA line. Structure and function claims are permitted; disease and treatment claims are not. A single non-compliant statement is a regulatory exposure.

HOW AN AI AGENT HANDLES IT

A clinical and technical assistant answers product and protocol questions drawn from compliance-reviewed product information, protocol guidance, and clinical references, with the DSHEA structure-versus-function boundary enforced. It explicitly does not generate disease or treatment claims and instead routes those questions to a designated compliance reviewer. A companion protocol builder helps licensed practitioners convert diagnoses and clinical goals into personalized, DSHEA-compliant supplement protocols and route the order into fulfillment, again surfacing only structure- and function-supported content and escalating anything that crosses the line. The same guardrails wrap competitor analysis, playbook Q&A, and rep training.

WHY IT MATTERS

It gives the commercial and practitioner-facing teams speed and consistency while making the compliance boundary a built-in property of the system rather than something a rep has to remember in the moment.

BUILT FOR	Commercial, Customer Success, Practitioner Enablement, Compliance
KEY KPIs	Response time, claim-compliance rate, escalation accuracy
TYPICAL IMPACT	Fast, consistent, on-label guidance with non-compliant claims routed to a reviewer by design

4 Research, Contract, and Utilization Analytics

THE PROBLEM

In a CRO, the answers to important questions are buried in systems only a few people can query: project, timesheet, and utilization data across Financial Force, OpenAir, Workday, and Salesforce; executed MSAs and study contracts; and multi-million-dollar syndicated datasets that route through a small BI or analyst team.

HOW AN AI AGENT HANDLES IT

A natural-language analytics agent lets PMs, ops, and finance ask questions across project and utilization data ("what is utilization on Study X this month?", "which projects are trending over budget?", "who has capacity in Q3 in Therapeutic Area Y?") and returns answers with drill-down, encoding the CRO's own business definitions so users do not need to know the schema. A companion legal-search agent searches executed MSAs, statements of work, and change orders in natural language ("have we agreed to this payment-term language with sponsor Y?", "which MSAs allow subcontracting without prior consent?") and returns ranked passages with the source document and section. The same shape extends to CV search and syndicated-dataset querying, so strategists can interrogate panel data directly.

WHY IT MATTERS

It turns multi-day analyst loops into minutes and makes expensive datasets and executed contracts usable directly by the people who need the answers, rather than routed through a bottleneck.

BUILT FOR	Project Management, Finance, Legal, Business Development, Research Strategy
KEY KPIs	Time-to-answer, analyst-loop reduction, dataset and contract accessibility
TYPICAL IMPACT	Multi-day loops become minutes, self-serve access to project, contract, and research data

5 Commercial Account Intelligence

THE PROBLEM

Preparing a business or quarterly review, or responding to an RFP, means hours of manual data assembly per account and inconsistent output across the commercial team.

HOW AN AI AGENT HANDLES IT

A review generator pulls the account's purchase history, product mix, year-over-year trend, e-commerce and ERP data, support history, and contract status, and produces a standardized business or quarterly review the rep takes into the meeting. An RFP and proposal agent decomposes a new RFP, pulls matching case studies and prior submission language from the proposal knowledge base, runs targeted external research on the prospect, and drafts a first-pass response in house style with citations, with an evaluator scoring each output against a rubric so quality stays consistent across consultants.

WHY IT MATTERS

It replaces hours of manual assembly with consistent, standardized output, so leadership can compare account health and proposal quality across the whole team.

BUILT FOR	Commercial, Sales, Account Management, Proposals
KEY KPIs	Review and proposal prep time, output consistency, win rate
TYPICAL IMPACT	Hours of assembly removed, standardized reviews and proposals, consistent quality across the team

05 — COMMON QUESTIONS, ANSWERED

Common questions, answered

The objections here are about regulatory survival, IP, and model control. Here are straight answers.

Q1 **“Anything touching complaints or quality has to survive an FDA inspection. Does yours?”**

That is the design intent. The complaints workflow produces a structured, consistent record with a full audit trail: every classification and extraction is logged with its inputs, source, output, reviewer, and timestamp, and a human approves before a record is finalized. Consistent, defensible records are exactly where manual intake tends to fail an audit, and where a governed agent strengthens your position rather than weakening it.

Q2 **“How do you keep supplement content on the right side of DSHEA?”**

The structure-versus-function boundary is built into the agent, not left to a rep's judgment. It grounds answers in compliance-reviewed sources, generates only structure- and function-supported content, and routes any question that crosses into disease or treatment territory to a designated compliance reviewer. The same guardrails apply to protocol building, competitor analysis, and training content.

Q3 **“Our clinical data, contracts, and IP cannot leak, and cannot train someone's model.”**

Deploy single-tenant so proprietary data never leaves your environment, and nothing you put into the platform trains any model or appears in logs. We hold BAAs and DPAs with model providers, and everything is encrypted. For CRO and device buyers this is the baseline, and it is why contract search, CV search, and clinical analytics run inside your own instance.

Q4 “We need to choose and validate our models, not be locked in.”

StackAI is model-agnostic with automatic fallback, and it includes an evaluator so you can run models head-to-head on your own criteria and document why an output is reliable. That is the posture a validation-minded, GxP-adjacent buyer needs: no vendor lock-in on the model layer, and a measurable basis for trusting an output.

Q5 “How is it priced, and how do we prove it before committing?”

A predictable monthly subscription, not usage-based, starting with a three-month proof of value on up to three workflows, all model tokens included, with an AI strategist and forward-deployed engineer in a weekly build-and-support session. We measure ROI on your own volumes (order cycle time, complaint intake hours, analyst-loop reduction) before any annual commitment.

06 COMPOUNDING VALUE

How these use cases compound

01 **Order operations and complaints intake** take the two heaviest regulated-operations loads off the team, one commercial, one quality, both with a defensible trail.

02 **Claims-compliant guidance** lets the commercial and practitioner side move fast without creating regulatory exposure.

03 **Research, contract, and utilization analytics** unlock the data and documents currently trapped behind a small analyst team.

04 **Commercial account intelligence** compounds all of it into faster, more consistent reviews and proposals.

Each runs inside the same governed, single-tenant, model-agnostic platform, so the regulatory posture is consistent across every workflow.

07 — THE $ VALUE

The value case

METRIC	BEFORE → AFTER AGENTS	IMPACT
Case-to-order cycle	before — **Hours of manual rekeying** after — **Seconds on routine cases**	Faster fulfillment, fewer corrections
Complaint intake	before — **Slow, inconsistent** after — **Fast, structured, defensible**	Reduced FDA audit risk
Claim compliance	before — **Depends on rep judgment** after — **Guardrailed by design**	Reduced regulatory exposure
Analyst loops	before — **Multi-day** after — **Minutes, self-serve**	Data and contracts unlocked
Review and proposal prep	before — **Hours per account** after — **Standardized, consistent**	Commercial time returned

08 WHY STACKAI

Why StackAI for life sciences

StackAI is the AI transformation platform for healthcare: a HIPAA- and GDPR-compliant operating system for enterprise agentic AI, built by MIT PhDs and deployed across more than 200 organizations in regulated industries.

Regulatory defensibility

Full audit trail on quality-relevant actions, human-in-the-loop, and consistent structured records for FDA readiness.

Claim guardrails

Built-in DSHEA structure-versus-function enforcement with escalation to a human reviewer.

IP protection

Single-tenant deployment, no model training on your data, BAAs and DPAs with providers.

Model-agnostic and validatable

No LLM lock-in, automatic fallback, and a built-in evaluator for head-to-head model comparison.

Enterprise reach

500+ integrations across Salesforce, Rootstock and other ERPs, SharePoint, Snowflake, Workday, and QMS platforms.

No-code builder

Quality, regulatory, and commercial experts build and change agents without an engineering ticket.

09 THE BOTTOM LINE

The bottom line

The operational and commercial work that consumes life-sciences and medtech teams is repetitive, documentation-heavy, and regulated at every step. Governed AI agents, built on approved sources, guardrailed on claims, defensible under audit, and running where your IP stays yours, take that load off the team while strengthening the regulatory position rather than putting it at risk.

Want to see this on your workflows?

Book a working session with the StackAI healthcare and life-sciences team to scope a single-tenant proof of value on your heaviest regulated-operations load, with ROI measured on your own numbers.