

**AKTANA**

**Intelligence Briefing**

# **Agentic AI: What Life Sciences Leaders Need to Know**

A practical reference for commercial life sciences teams exploring the next evolution of AI-driven customer engagement.

# Table of Contents



## The Future of AI in Life Sciences

An Introduction from John Vitalie

03



## Humans Must Be in the Loop, Not on the Loop

A Reflection from Derek Choy

05

Agentic AI & Human Oversight

07



## Transparency and Traceability as Foundations of Trust

Insights from Derek Choy

17

Agentic AI & Transparency

18



## Context-Aware AI for Effective Decision Making

A View from Derek Choy

28

Agentic AI & Continuous Improvement

29



## Practicality and Integration as Cornerstones of Future-Ready AI

A Conclusion from John Vitalie and Derek Choy

37



# The Future of AI in Life Sciences

## An Introduction from John Vitalie



**John Vitalie**, CEO, Aktana

“ To our friends and colleagues in Life Sciences,

We're at a pivotal moment in our industry. Rising cost pressures and global volatility are creating a deeper urgency to maximize investment returns. As intelligent automation programs mature beyond the early stage, expectations are also rising to rebalance content budgets. To ensure every action delivers economic value, organizations must navigate regulatory demands while also optimizing resources in real time. Achieving a lasting competitive advantage should be the definitive strategic goal when implementing practical agentic AI solutions.

AI has been reshaping our industry for years, transforming how we engage with healthcare professionals, optimize workflows, and make data-driven decisions. The rapid progress is exciting, but it also raises important questions:

Can (and should) AI move beyond recommendations to truly drive action autonomously?

How do we keep AI operations transparent, compliant, and aligned with strategic goals?

We believe that embracing AI doesn't mean sidelining the very human expertise, collective experience, and individual judgment that drive results across our industry. Our vision is clear:

AI should empower your teams, augment their capabilities, align objectives to execution, and elevate human decision-making rather than replace it.

This ebook continues that conversation. Throughout the following pages, you'll find a practical perspective on how agentic AI can deliver measurable value today, while laying the foundation for a more autonomous, insight-driven future. With contributions shaped by years of real-world experience, we hope this resource provides clarity, context, and inspiration as you chart your agentic AI journey.”

Best,

John Vitalie, CEO, Aktana

**AKTANA**

# **Humans Must Be in the Loop, Not on the Loop**

Agentic AI & Human Oversight



# Humans Must be In-the-Loop, Not On-the-Loop

## A Reflection from Derek Choy

“ Reflecting on my decade of pioneering AI in life sciences, I’m reminded of the many challenges we faced along the way. When we first introduced AI-driven recommendations, we realized its success depended on more than just intelligence. Initial resistance from field teams, stemming from concerns over loss of control and skepticism toward unclear recommendations, highlighted a critical truth: **AI solutions must empower, not replace, human judgment.**

Today, I’m proud to see this philosophy drive Aktana’s Agentic AI solution. Recent months (and weeks) have brought tremendous advances in AI, and major industry announcements such as Veeva AI, Salesforce Agentforce’s support for MCP, and the Life Sciences Cloud digital labor and data acceleration partnerships. With growing willingness to embrace intelligent automation, the moment is finally right for the industry to realize the vision we set out a decade ago - and unlock its full value.

This vision is that Agentic AI has the potential to transform life sciences by streamlining operations and orchestrating omnichannel engagement, freeing professionals to focus more on strategy and outcomes. Consider this scenario:

- A **Brand Manager** reviews AI-suggested tactics based on a domain-specific knowledge base, confident each action is transparent and strategically-aligned.
- A **Commercial Ops lead** confirms that AI recommendations respect resource allocation brand tradeoffs and compliance rules, with safeguards for accuracy and traceability.
- A **Sales Rep** receives personalized HCP outreach suggestions - including compliant content - which she can quickly adjust to maintain her personal touch.
- A **Medical Science Liaison** prepares for HCP conversations with AI-curated clinical insights, focusing on what matters most.
- And of course the **HCP** benefits - accessing compliant information via field reps or even virtual agents, and easily scheduling follow-ups.

Notice that in all these scenarios, one principle holds: **AI doesn’t replace people - it elevates them.** None of this is about AI taking control away from the humans who make this all happen. **It’s about amplifying human potential while reinforcing control, trust, and compliance.**

Expert guidance is essential to ensure safeguards in agentic AI, especially as generative AI's potential for errors ("hallucinations") requires continuous human feedback. **The life sciences industry requires a constant equilibrium between innovation and compliance.** These principles have always been at the core of Aktana's collaborative AI approach.

While the list of tasks AI will handle independently continues to grow, AI should support—not dictate—decisions. When humans are merely **"on the loop"**, they're relegated to overseeing decisions without influence. Instead, they must be **"in the loop"**: central actors with AI augmenting their expertise, not marginalizing it.

At Aktana, we believe this is a defining moment. **Agentic AI could cut down commercial planning and execution timelines from up to 18 months to as little as 6 months. Those kinds of gains - whether in faster product launches, smarter resource use, or more personalized HCP engagement - are transformative.** But they're only achievable if AI is trusted and adopted by your teams. This, in turn, requires AI solutions that are intuitive, explainable, and clearly aligned with strategic goals - all hallmarks of a human-centered AI approach.”



**Derek Choy**  
Co-founder and Chief Product  
Officer, Aktana

Frequently Asked Questions About

# Agentic AI & Human Oversight

Click on a question to jump to answer

- Q1** What is Agentic AI, and how does it relate to conventional AI approaches used in life sciences today?

---

- Q2** What are Agents, and how do they contribute to decision-making in complex systems?

---

- Q3** What's the difference between Agentic AI and AI agents?

---

- Q4** What's the difference between single-agent vs. multi-agent in AI, and why does it matter?

---

- Q5** I already have a CRM and Marketing Automation system in place. How does an AI-powered agentic system enhance what I'm already using?

---

- Q6** Can I personalize HCP engagement at scale without overwhelming my team with more data and complexity?

---

- Q7** What's the difference between a domain-specific Agentic AI system and a general AI model? Why does it matter for BioPharma?

---

- Q8** How does Agentic AI integrate with existing life sciences systems like CRMs and Marketing Automation platforms?

---

- Q9** How can AI-driven agents adapt in real-time to human input?

## Question 1

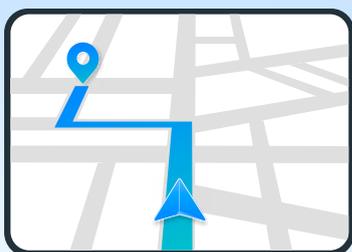
# What is Agentic AI, and how does it relate to conventional AI approaches used in life sciences today?

Agentic AI goes beyond analyzing data and providing recommendations - it takes action. It **plans, orchestrates, and executes tasks autonomously**, adapting to real-time changes. With the introduction of Model Context Protocol (MCP), it can also plug into enterprise tools and data. Unlike traditional AI, which requires defined data, predefined rules or algorithms, and requires human intervention, **Agentic AI dynamically optimizes workflows while maintaining human oversight.**

A useful analogy is the difference between GPS navigation and a self-driving car. Conventional AI is like GPS - it suggests the best route but leaves execution to the driver. Agentic AI, like a self-driving car, adjusts to real-world conditions, ensuring continuous optimization.

While Agentic AI introduces automation, fully autonomous AI - especially LLM-driven agents - has limitations in regulated industries. Generative AI systems are prone to hallucinations (plausible-sounding but incorrect outputs) and can drift outside of approved content or data if not constrained. A structured approach ensures AI supports decision-making **without sacrificing reliability, transparency, or strategic alignment.**

### Conventional AI - Like a GPS



- Suggests route
- Human drives
- Static inputs
- Limited adaptability

### Agentic AI - Like a Self-Driving Car



- Plans and executes
- Adjusts to real-time changes
- Interacts with tools/data
- Maintains human oversight

## Question 2

# What are Agents, and how do they contribute to decision-making in complex systems?

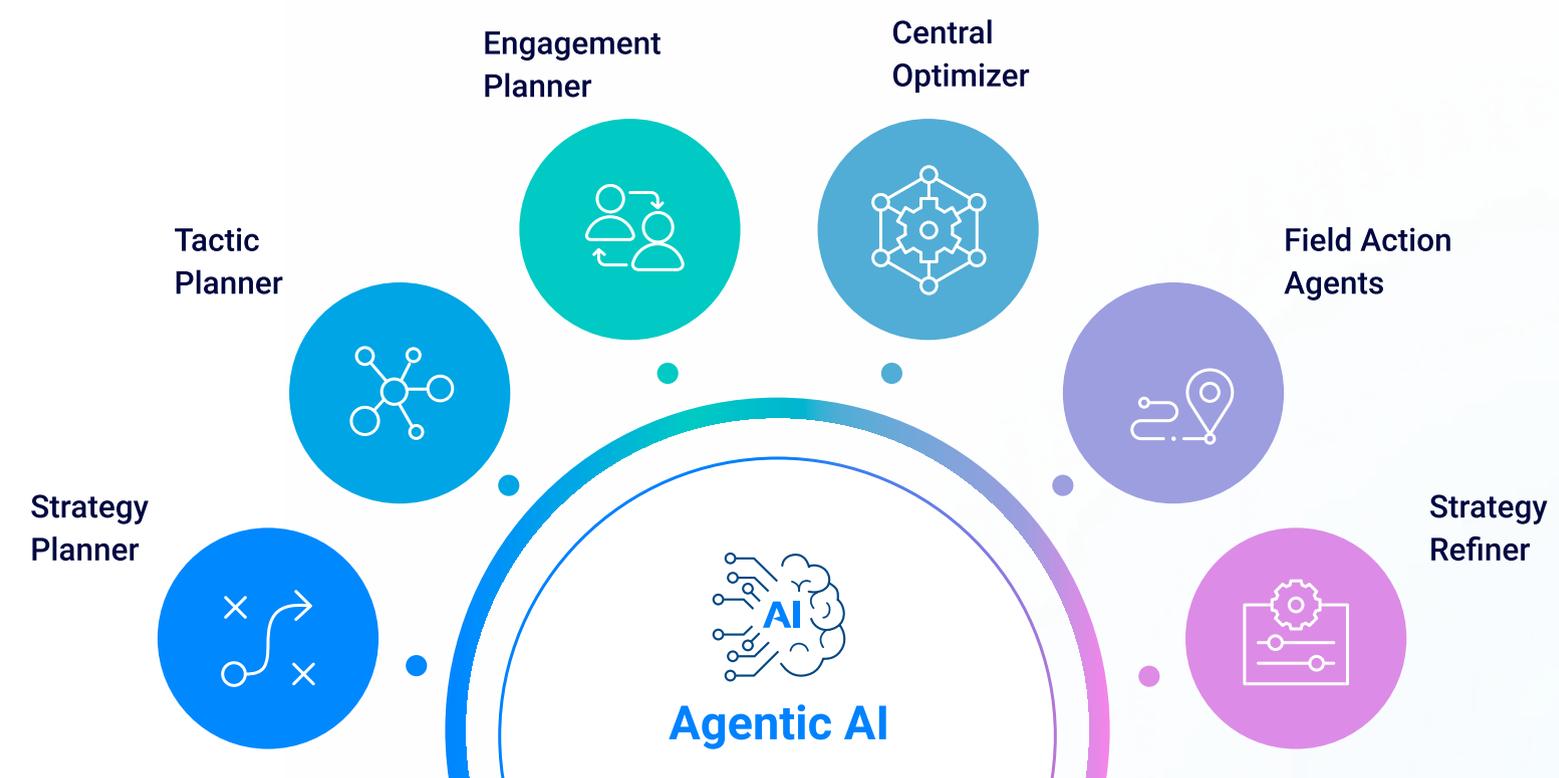
In Agentic AI systems, “agents” are like specialized domain experts working together toward a common goal. Each has a clearly defined role, aligned to a strategic objective, and contributes to the overall decision-making process.

Think of Agentic AI as a team of intelligent specialists, modeled on how life sciences teams already operate:

- **Strategy Planners** define brand-level goals and KPIs to guide downstream decisions.
- **Tactic Planners** translate strategy into personalized HCP-level actions - what to say, when, and through which channel.
- **Engagement Planners** coordinate across brands and teams to align resources and avoid overlap.
- **Central Optimizers**, most critically, work with a knowledge base and orchestrators to sequence actions based on value, constraints, and compliance rules.
- **Field Action Agents** deliver timely suggestions and insights to reps and MSLs to support execution.
- **Strategy Refiners** track outcomes and KPIs to improve future recommendations.

This distributed model makes Agentic AI systems more adaptable, scalable, and transparent than single-model AI approaches - especially in the high-stakes, fast-moving environment of life sciences.

As the ecosystem evolves, standards like the **Model Context Protocol (MCP)** and **Agent-to-Agent (A2A)** communication are emerging to enable seamless interaction between agents and enterprise tools. These innovations will make it easier to extend capabilities, improve interoperability, and accelerate the adoption of Agentic AI across teams and technologies.



### Question 3

## What's the difference between Agentic AI and AI agents?

While the terms sound similar, there's a key difference between Agentic AI and AI agents.

- **AI agents** are individual specialists. Each agent performs a specific task. In life sciences, an AI agent might analyze patient data for clinical trials or generate email content for HCP engagement. These agents operate within a limited scope and often rely on external triggers or human intervention to take action.
- **Agentic AI** is the entire system - a coordinated network of AI agents working together to plan, execute, and optimize complex workflows. For example, in a drug launch, Agentic AI **orchestrates multiple AI agents** to align engagement strategies, ensure compliance, and adapt to real-time market dynamics.

Agentic AI represents a **higher level of coordination and adaptability**, allowing AI-driven processes to function more dynamically than individual AI agents working in isolation.

### Question 4

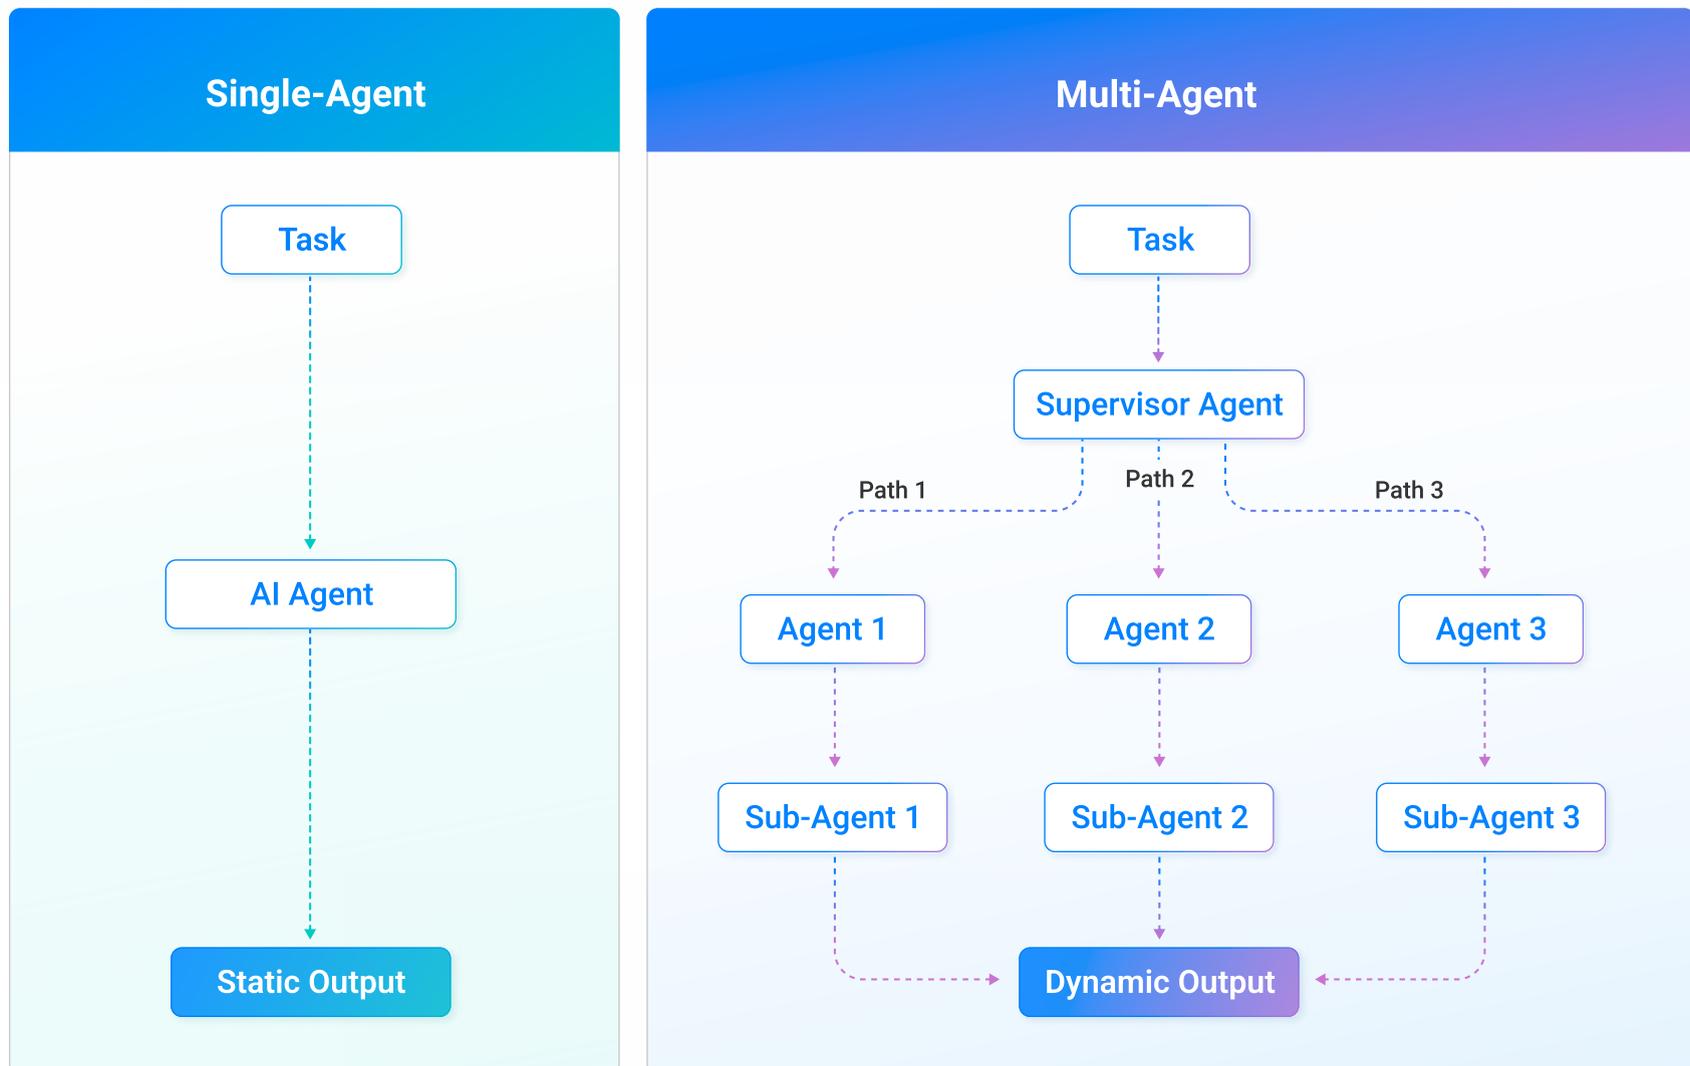
## What's the difference between single-agent vs. multi-agent in AI, and why does it matter?

AI agents are intelligent entities that can make decisions and take action, **but how they work together determines their effectiveness.**

- **Single-agent AI** operates independently, like a solo expert tackling a project. While effective for specific tasks, a single agent **lacks the ability to coordinate across complex workflows** and may struggle with dynamic, multi-step decision-making.
- **Multi-agent AI** functions as a team of specialists, with different agents handling distinct responsibilities. They **communicate, share information, and collaborate** to solve problems in parallel, leading to more efficient, adaptable, and scalable decision-making.

Multi-agent AI enables **specialization, coordination, and adaptability**, making it ideal for complex, regulated industries. In life sciences, for example, one agent can monitor compliance, another can optimize omnichannel engagement, and others can adjust field execution based on real-time data. If one agent encounters an issue, others can compensate, ensuring continuity and resilience.

By distributing intelligence across multiple AI agents, organizations gain a system that is **more flexible, scalable, and capable of handling real-world complexities** without sacrificing compliance or efficiency.



## Question 5

I already have a CRM and Marketing Automation system in place. How does an AI-powered agentic system enhance what I’m already using?

While CRM and Marketing Automation systems remain essential for managing HCP engagement, there is room for greater efficiency in how they are used. At present, manual intervention is often required for decision-making, adjusting strategies, and resolving errors. An AI-powered agentic system enhances these platforms by making them **more dynamic, responsive, and intelligent**, without disrupting existing workflows. This is made possible through a few core capabilities that define how agentic AI systems enhance CRM and marketing automation:

- **Autonomous Planning** – AI determines the best sequence of steps based on goals, real-time HCP behavior, and market conditions rather than relying on static workflows.

- **Adaptive Execution** – When prescription trends, regulatory requirements, or engagement patterns shift, AI adjusts strategies automatically to maintain relevance and compliance.
- **Tool Selection and Use** – AI agents interact with CRM, content platforms, and analytics tools as needed, reducing integration complexity and ensuring seamless execution.
- **Memory and Context Management** – AI retains engagement history across channels, ensuring continuity. If an HCP moves from digital to in-person interactions, AI adapts the outreach accordingly.
- **Feedback Integration** – AI refines strategies in real time, prioritizing tactics based on engagement success. If oncologists engage more with a specific content format, AI optimizes recommendations accordingly.

**For life sciences companies, the goal isn't just automation - it's making sure AI-driven engagement works within real-world workflows, aligns with strategic priorities, and meets regulatory standards.** A multi-agent approach with human oversight can support this by keeping AI adaptable, transparent, and trusted by teams.

## Question 6

# Can I personalize HCP engagement at scale without overwhelming my team with more data and complexity?

Scaling HCP engagement is like planning a mountain climb. Traditional engagement strategies follow a fixed path based on historical data, but the real world doesn't work that way. Competitor activity shifts, regulations evolve, and HCP behaviors change. **Without real-time adjustments, teams either stick to outdated plans or scramble to adapt manually.**

Just like an experienced guide adjusts a route based on conditions, Agentic AI continuously refines engagement strategies in real time. Instead of requiring teams to manually reassess every step, more advanced AI optimizes **outreach, messaging, and engagement tactics** to stay aligned with market realities.

But even the best AI-driven systems aren't meant to navigate the path alone.

**A human-in-the-loop approach ensures that AI adapts to changing conditions while keeping commercial teams in control.** Rather than replacing human judgment, this approach ensures teams can act on AI-driven recommendations with confidence, keeping personalization scalable and manageable.

## Question 7

# What's the difference between a domain-specific Agentic AI system and a general AI model? Why does it matter for BioPharma?

General AI models are broad and adaptable, but they lack the deep industry knowledge, regulatory awareness, and real-time adaptability required in life sciences. They provide insights but rely on human intervention for execution and must be retrained to incorporate new data.

Domain-specific Agentic AI is built for complex, highly regulated environments like biopharma. It integrates compliance requirements, understands industry-specific challenges, and continuously adapts without retraining. By breaking down complex problems into smaller, testable steps, it reduces errors, ensures accuracy, and enhances decision-making, critical for regulatory adherence and commercial execution.

In life sciences, where precision and compliance are essential, AI must do more than automate tasks. It needs to **support decision-making in a way that is transparent, traceable, adaptable, and aligned with business objectives.**

Aspect	General AI Model	Domain-Specific Agentic AI
Regulatory Awareness	No built-in compliance checks	Designed with industry regulations in mind
Adaptability	Requires retraining for updates	Adjusts in real time based on new data
Actionability	Provides insights but requires human execution	Automates workflows while keeping teams in control
LLM Accuracy	Can generate errors or hallucinations	Breaks down problems into testable steps to ensure accuracy
Workflow Integration	Works in isolation, requiring manual setup	Seamlessly connects with life sciences systems

## Question 8

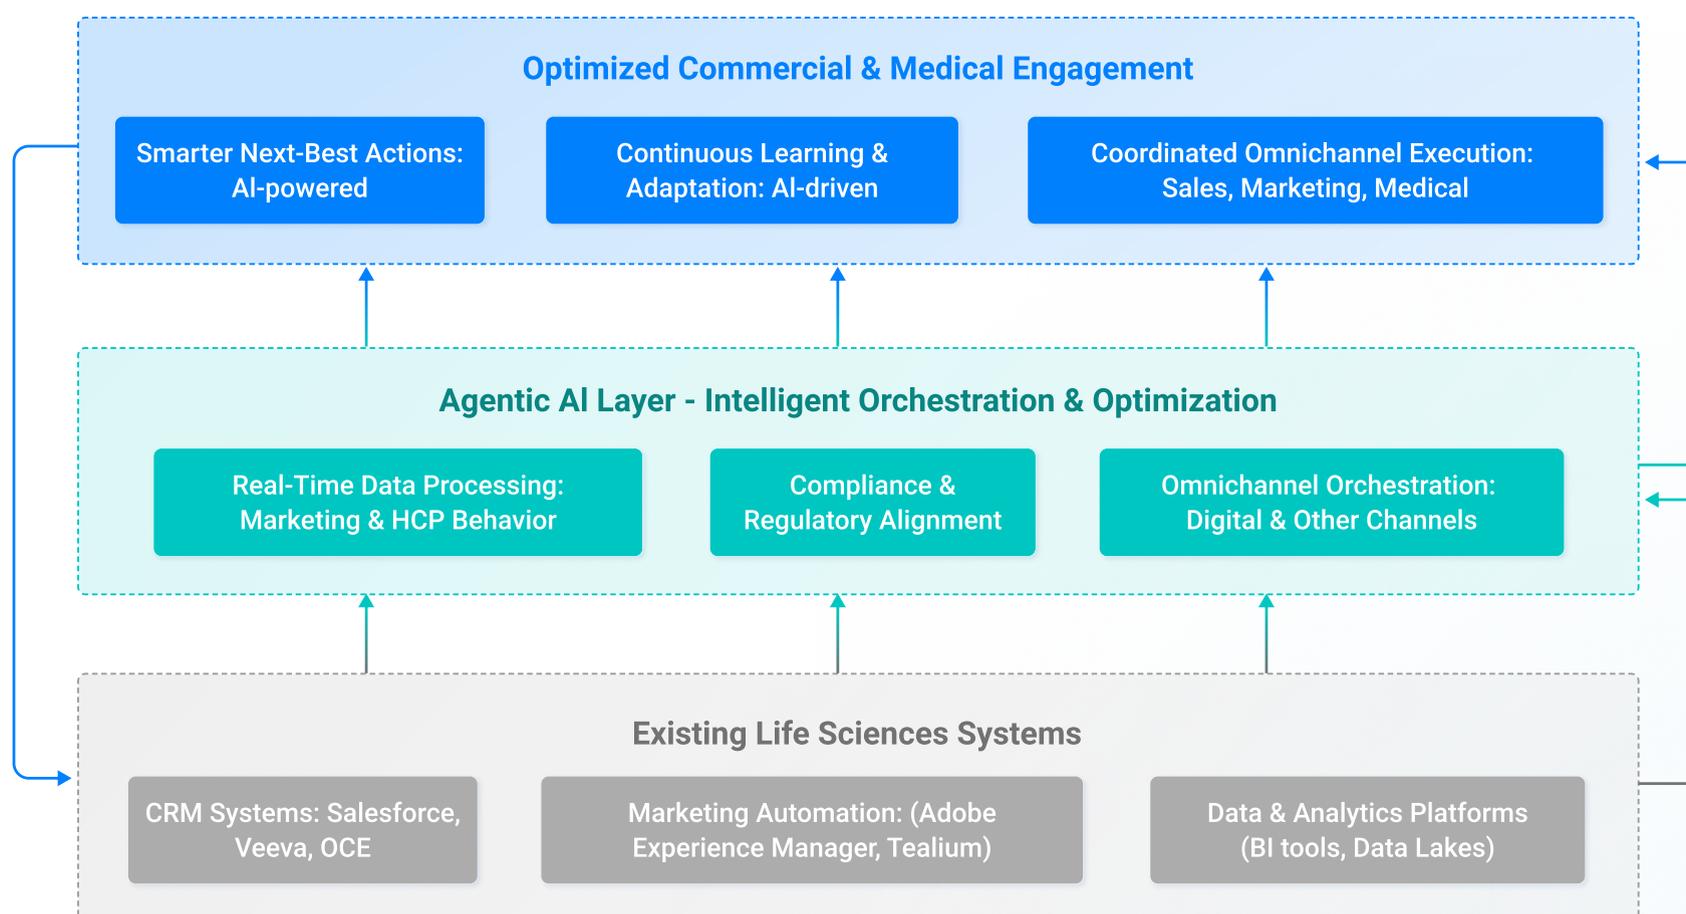
# How does Agentic AI integrate with existing life sciences systems like CRMs and Marketing Automation platforms?

Life sciences organizations rely on CRMs, Marketing Automation, and analytics platforms to manage HCP engagement and commercial execution. **The challenge is that these systems often operate in silos, requiring manual coordination and struggling to adapt to real-world complexities.** Agentic AI enhances these tools by working as an intelligent layer on top of existing infrastructure, enabling real-time orchestration and adaptive decision-making.

Instead of following static workflows, AI agents **analyze live data, market conditions, and HCP behavior to refine engagement strategies on the fly.** It ensures that sales reps, MSLS, and marketing teams receive coordinated, context-aware recommendations, improving efficiency without increasing complexity.

This orchestration isn't tied to a single system. Flexible integration ensures AI-driven decision-making works across platforms, rather than being restricted to a specific CRM or marketing automation tool. Companies like Aktana facilitate this through **solutions that allow AI to interact with multiple systems** - including Salesforce, Veeva, Adobe Experience Manager, and OCE - so organizations can apply AI-driven insights **within their existing workflows.**

By embedding intelligence into the omnichannel ecosystem, **Agentic AI transforms CRMs and marketing automation from static execution tools into dynamic, adaptive platforms.** This allows commercial and medical teams to scale engagement efficiently, ensuring that every action is optimized for real-time market conditions.



## Question 9

# How can AI-driven agents adapt in real-time to human input?

AI-driven agents in life sciences operate within a **human-in-the-loop framework, balancing automation with oversight** to ensure adaptability, compliance, and strategic alignment. Rather than replacing human decision-making, **Agentic AI refines engagement strategies in real time, adjusting based on new data and expert input.**

This adaptability is driven by key capabilities, including:

- **Real-Time Adjustments** – Teams can override AI-driven recommendations to stay aligned with market conditions and compliance needs.
- **Context Retention** – AI learns from human input, refining future recommendations without repeating past mistakes.
- **Adaptive Workflows** – AI integrates feedback continuously, adjusting strategies based on **HCP behavior, rep interactions, and regulatory updates.**
- **Strategic Oversight** – AI supports, rather than replaces, human decision-making, ensuring engagement remains effective and compliant.

But these alone aren't enough. AI needs **both structure and context** to be effective in life sciences. Without these, black-box AI risks making opaque, misaligned decisions that introduce inefficiencies or compliance risks.

- **AI needs a central intelligence layer to orchestrate decisions** – To align execution with business goals, AI must continuously weigh competing priorities, optimize actions across channels, and integrate constraints like compliance and resource availability.
- **AI must be guided by contextual knowledge** – Industry expertise, historical data, and regulatory frameworks provide the foundation AI needs to make informed, transparent, and strategically sound decisions.

Aktana enables this by ensuring AI-driven execution is **not just automated, but adaptive, explainable, and aligned with real-world business needs.**

**AKTANA**

# **Transparency and Traceability as Foundations of Trust**

Agentic AI & Transparency



# Transparency Is Non-Negotiable in AI: Derek Choy Explains Why

A pivotal lesson from our early AI implementations is that transparency is non-negotiable. In one of our earliest deployments with a major pharma company, we observed firsthand how skepticism arose when AI recommendations felt like unexplained commands rather than thoughtful suggestions. Field representatives questioned the value of AI-driven recommendations because they lacked clear, relatable reasoning. This was understandable - after all, these representatives had significant experience and relied on their judgment and deep customer relationships. **For them, trust in AI-driven recommendations was driven not just by accuracy but also intuitive, relatable context that “felt right” in the light of their own professional insights.**

The turning point came when we started providing transparent, contextual reasons based on relevant, real-world data for every AI recommendation. This included context like recent HCP interactions, historical prescribing patterns, peer comparison insights, and even subtle indicators like the HCP’s preferred communication channel and timing. **This made AI recommendations feel actionable and logical to reps, increasing adoption and confidence in the system.**

Transparency is also crucial for strategic brand management. When brand leaders can clearly trace the reasoning behind AI-driven strategies, they can more confidently delegate strategic decisions to semi-automated or even fully-automated systems. This enables brand managers to quickly validate, adjust, or embrace recommended strategies without extensive manual review.

“ Today, Aktana’s **Agentic AI** approach institutionalizes this principle of transparency. Every recommendation our system provides is **fully traceable and understandable**, ensuring compliance with strict regulatory standards and also meaningful human oversight.”

**Derek Choy,**  
Co-founder & Chief Product Officer, Aktana

Our central orchestration engine serves as a transparency hub, providing clear visibility into every decision: why it was recommended, what data informed it, and how it aligns with strategic goals. This explainability is vital for compliance teams, ensuring complete auditability to satisfy stringent regulatory requirements.

Ultimately, transparency and traceability are essential for both building trust and also maintaining sustained confidence and adoption. By embedding these qualities into our Agentic AI platform, we empower decision-makers with clarity, control, and trust - paving the way for the scalable, reliable omnichannel orchestration essential to the success of today’s pharma organizations.

Frequently Asked Questions About

# Agentic AI & Transparency

Click on a question to jump to answer

**Q1** What are the potential benefits of well-designed autonomous agents in life sciences?

---

**Q2** How does Agentic AI differ in functionality and application from Generative AI?

---

**Q3** What are the risks of “black box” AI in life sciences, and how does Agentic AI mitigate them?

---

**Q4** Why is explainability critical for AI adoption in commercial life sciences teams?

---

**Q5** Regulatory compliance is a huge concern in life sciences. How can AI-driven agents ensure compliance while still enabling innovation?

---

**Q6** Why is advanced orchestration still essential in the age of Agentic AI?

---

**Q7** How can Agentic AI streamline operations and team workflows in life sciences?

## Question 1

# What are the potential benefits of well-designed autonomous agents in life sciences?

In a regulated, data-intensive industry like life sciences, autonomous AI agents have the potential to significantly improve how organizations engage, operate, and scale - if designed and deployed with care. The following examples highlight the kinds of benefits that a well-architected Agentic AI system could deliver:

### 1. Dynamic Communication Planning and Personalization

Agents can now leverage AI to craft and deliver highly personalized and contextually resonant communication plans. They can dynamically adapt messages, channels, and timing in real time based on detailed HCP behavior, brand objectives, and engagement patterns. This can result in communications that truly speak to each individual, maximizing impact, improving content utilization, and freeing up field teams to focus on deeper relationships.

### 2. Coordinated Execution Across Channels and Teams

Multiple agents can operate in parallel to orchestrate omnichannel strategies, monitor field capacity, adjust next-best actions, and enforce brand priorities. This, in turn, enables faster campaign execution and cross-functional alignment without sacrificing control.

### 3. Transparent, Auditable Decision-Making

Every recommendation from the system should come with clear reasoning - from content suggestions to tactical adjustments. This supports explainability, builds trust across teams, and ensures decisions are easy to validate and audit.

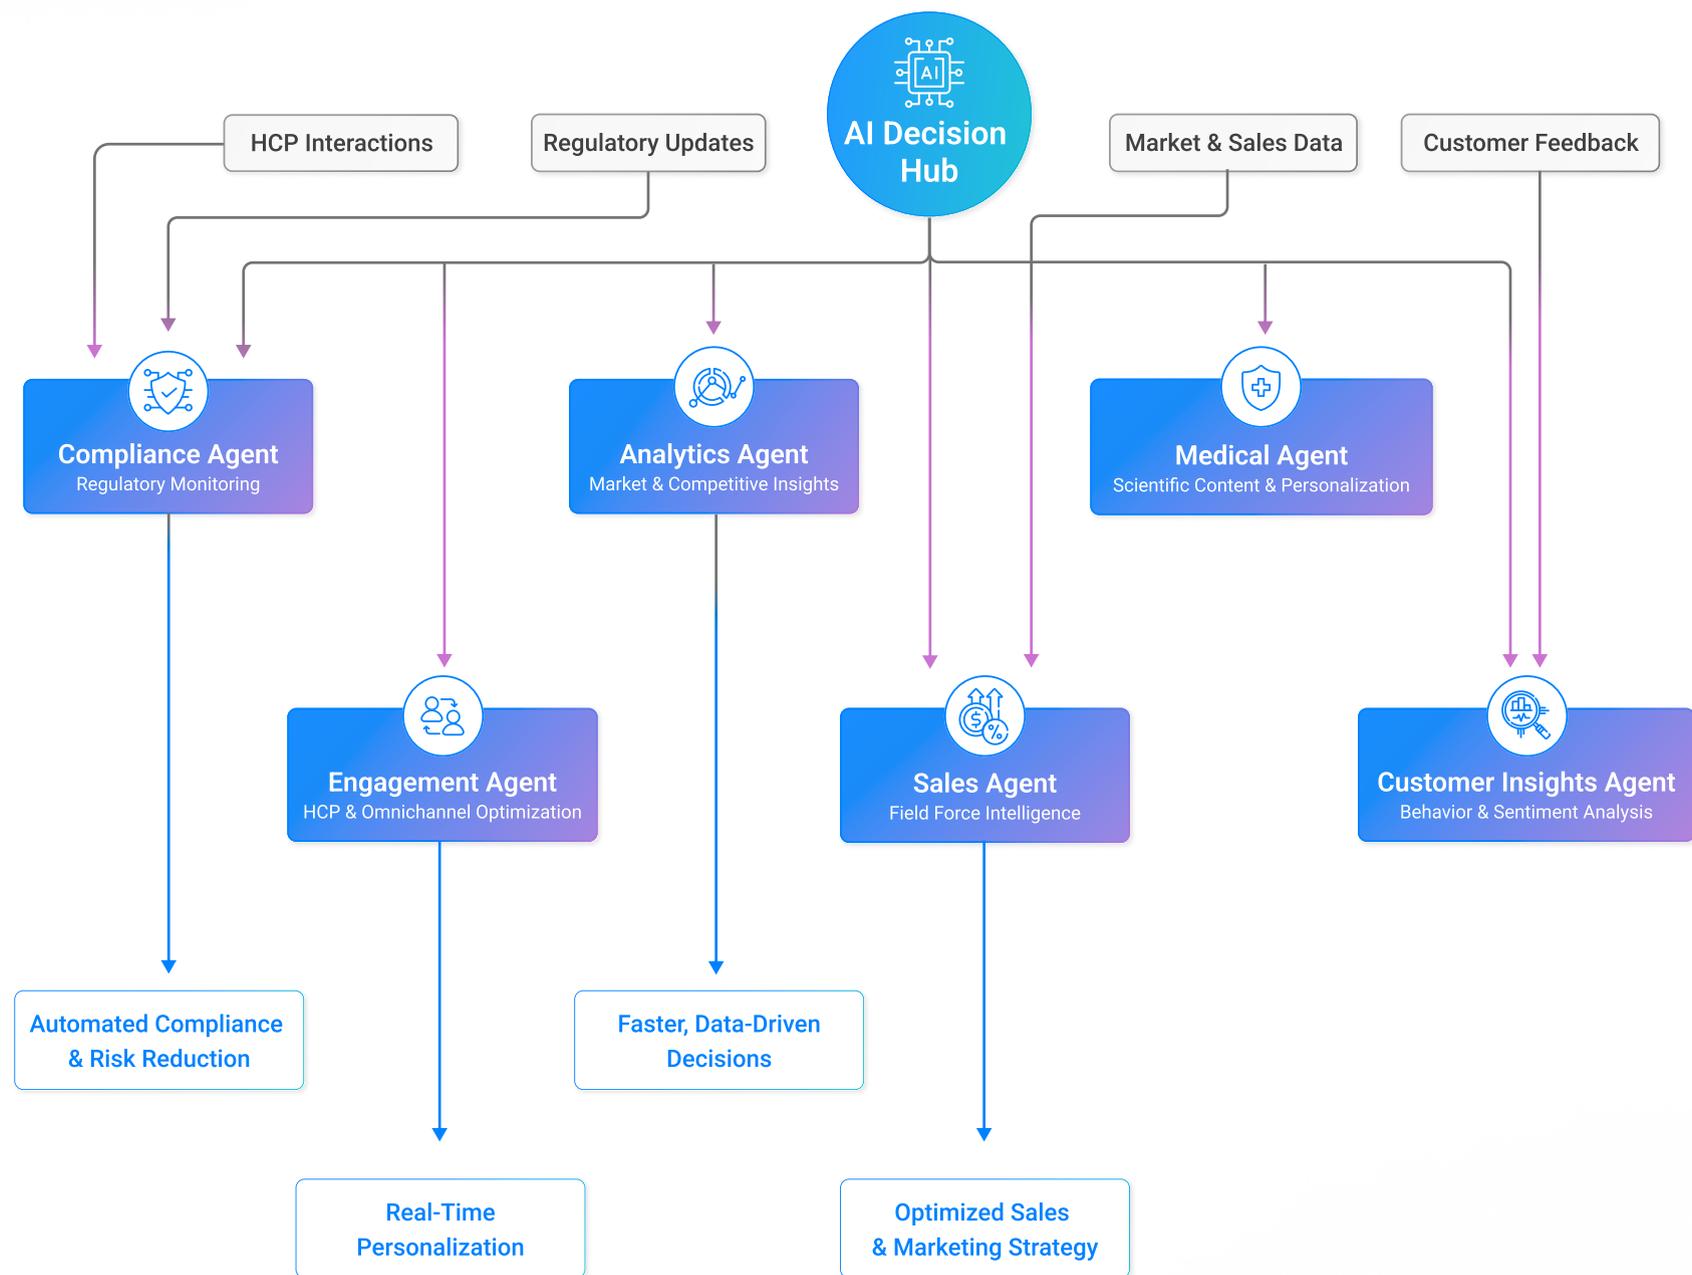
### 4. Adaptation Without Disruption

Agents will be able to autonomously respond to market and regulatory changes such as new product launches or label updates. By eliminating manual reconfiguration, an agentic system should remain agile even during periods of transition.

A multi-agent decision ecosystem enables AI agents to process data, apply expertise, and act autonomously in real time. This foundation for practical Agentic AI ensures transparency, human oversight, and alignment with strategy and regulations. Such architecture powers platforms like Aktana's, enabling scalable, trusted automation.

### How this works in practice:

This visual demonstrates the collaborative nature of an illustrative multi-agent decision ecosystem, showcasing how a wide range of specialist agents can work together to enable personalized communication, strategic execution, and transparent decision-making at scale.



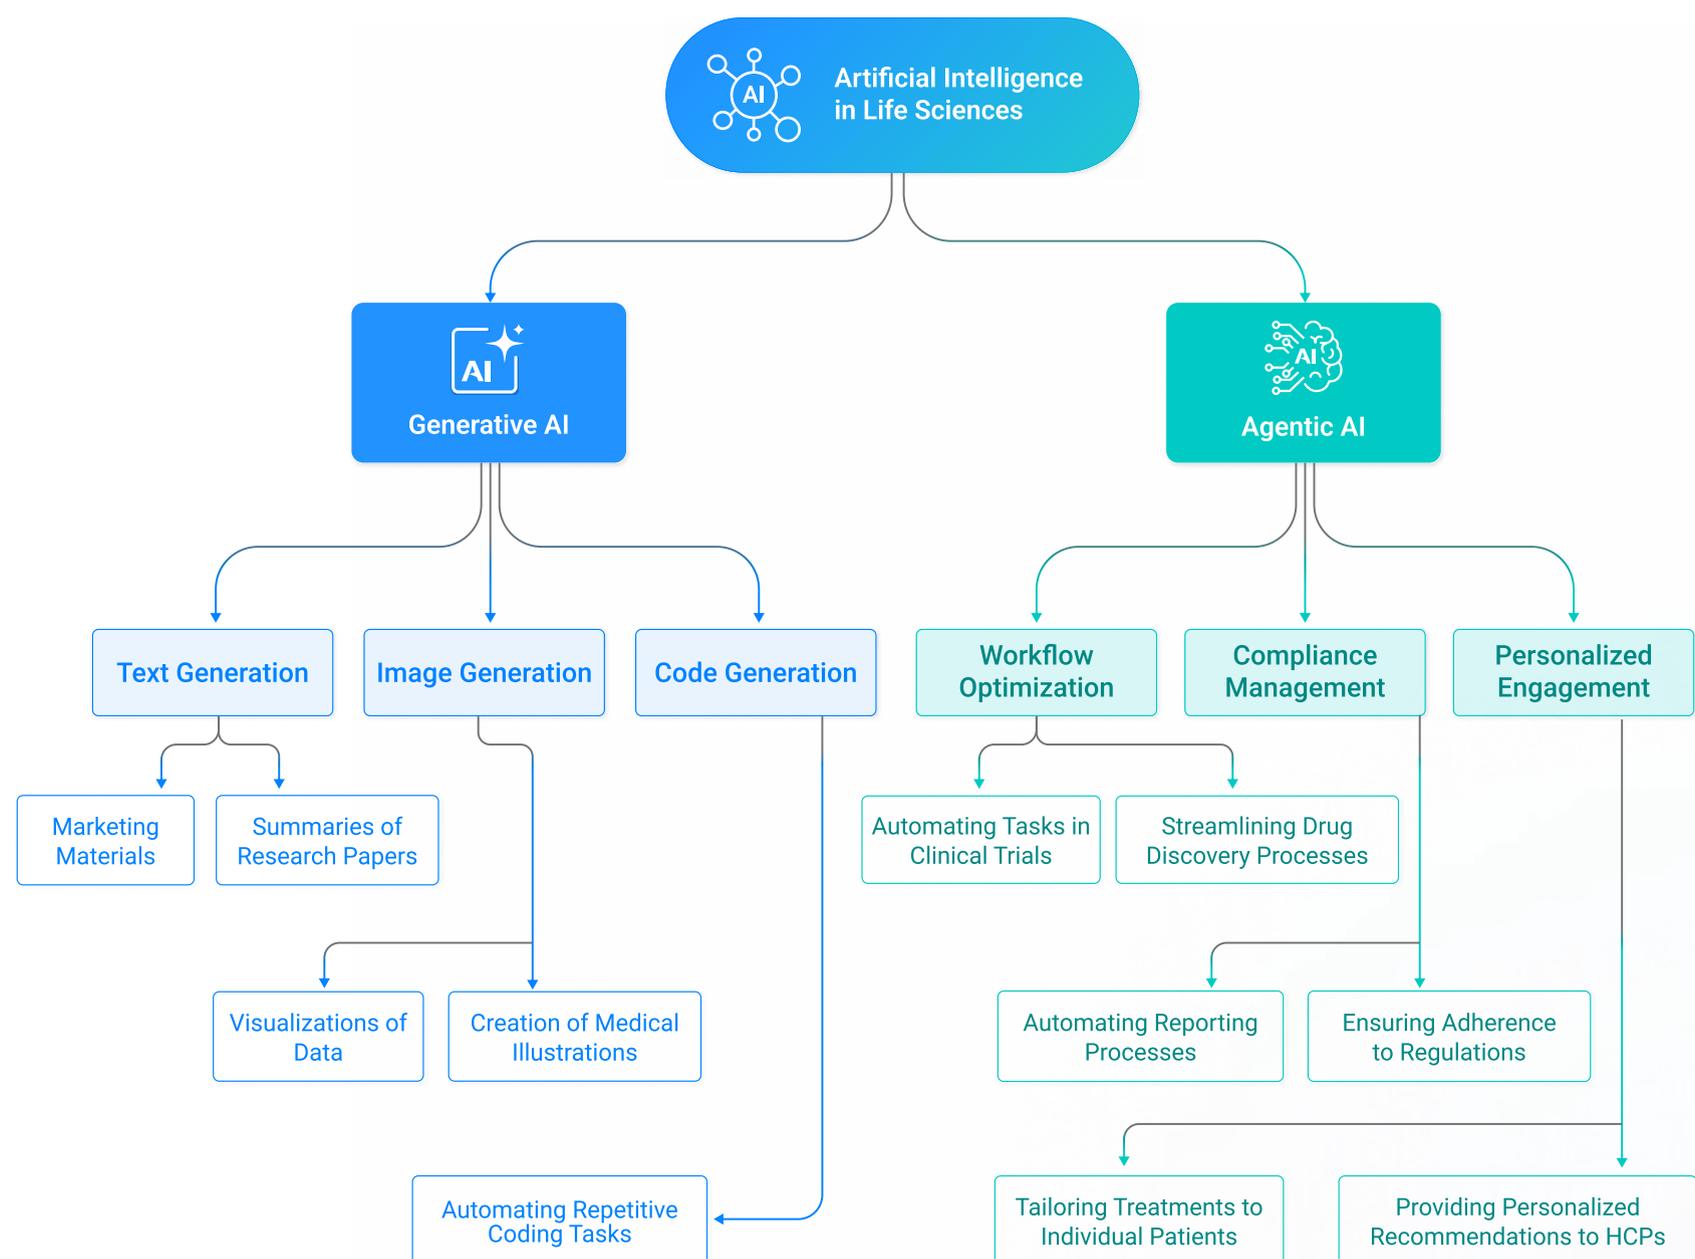
## Question 2

# How does Agentic AI differ in functionality and application from Generative AI?

While both Generative AI and Agentic AI advance artificial intelligence, they address distinct needs through different approaches.

- **Generative AI** focuses on content creation - generating text, images, or code - by identifying patterns in data. It automates content creation and summarizes information but does not make independent operational decisions.
- **Agentic AI**, on the other hand, expands the utility of Generative AI but is much more geared towards action and orchestration. It plans workflows, executes strategies in real-time, and coordinates tasks. It functions as a collaborative ecosystem of AI agents working together to optimize operations and ensure compliance.

In short, Generative AI enhances content production, while Agentic AI drives complex operations. Crucially, Agentic AI can utilize Generative AI outputs to craft highly personalized, nuanced communications, such as tailored emails, dynamic summaries, and context-specific content variations, and strategically deploy them in real time based on objectives and HCP responsiveness.



### Question 3

## What are the risks of “black box” AI in life sciences, and how does Agentic AI mitigate them?

Trust is paramount in life sciences, where decisions impact compliance, patient safety, brand integrity, and field adoption. “Black box” AI, which provides outputs without explaining their origins, presents significant risks:

- Field teams may dismiss recommendations they cannot understand.
- Brand leaders may struggle to validate AI-driven strategies.
- Regulatory teams may be unable to audit or defend decisions.

Well-designed and high-quality Agentic AI, if approached in the right way, can mitigate these risks with transparent, explainable reasoning, providing context such as HCP engagement data, prescribing patterns, and constraints for each recommendation. This is why separating out the agent which transparency orchestrates and optimizes is paramount. This inherent traceability allows brand teams to ensure strategic alignment, compliance officers to track a complete audit trail, and field users to trust suggestions as both accurate and logical.

### Question 4

## Why is explainability critical for AI adoption in commercial life sciences teams?

Explainability in AI means **understanding how and why an AI system makes a recommendation**, fostering trust and building confidence. ‘Black box’ AI breeds skepticism.

Here's why explainability is critical for AI adoption:

- **Builds Confidence:** Understanding the reasoning behind AI recommendations - the data, logic, and “why” - gives teams confidence to act, transforming AI into a trusted partner.
- **Validates Field Team Expertise:** Explainable AI aligns with field team experience, providing context that supports their judgment. This is particularly crucial when AI selects or sequences content. When reps and brand managers understand why a particular personalized message or communication arc was chosen - based on strategy, engagement, or behavioral signals - they are more likely to trust and use the system.
- **Ensures Accountability:** Transparency allows teams to validate, track, and ensure compliance. Understanding is key to accountability.
- **Facilitates Adoption:** When AI's decisions are clear, it is more likely to be used. Explainability removes barriers to adoption.

Ultimately, explainability empowers teams, making AI a trusted tool that turns data into actionable insights.

## Question 5

# Regulatory compliance is a huge concern in life sciences. How can AI-driven agents ensure compliance while still enabling innovation?

In life sciences, even a small error, like incorrect dosage information in AI-generated content, can trigger product recalls, regulatory investigations, and reputational damage. That is why, for Agentic AI to succeed in this space, **compliance cannot be an add-on; it must be foundational.**

From the outset, all AI-driven actions should be governed by clear business rules and subject to human oversight. These rules serve as guardrails, ensuring that even as agents operate autonomously, they never deviate from approved processes or content.

To further mitigate risk, **Agentic AI systems must undergo rigorous testing** - evaluated against a wide range of simulated scenarios, edge cases, and adversarial prompts. This ensures the system performs reliably under regulatory pressure, even in complex or unexpected conditions.

Transparency is equally essential. Like a detailed patient record, every AI decision must be **fully traceable**: what the system did, why it did it, and whether it stayed within regulatory boundaries. Tools like [MLflow in Databricks](#) help log and monitor these decisions, creating a defensible audit trail that supports compliance reviews and internal governance.

Embedding compliance, rigorous testing, and full transparency enables life sciences organizations to confidently innovate while meeting industry standards for speed and scale. This is especially important in content-driven workflows, where AI agents are responsible for not just delivering - but also selecting - appropriate, pre-approved messages and sequencing them within regulatory boundaries. Every step must be explainable, auditable, and compliant.

## Question 6

# Why is advanced orchestration still essential in the age of Agentic AI?

Agentic AI's strengths in data analysis, planning, and autonomous action make it ideal for near real-time applications like pre-call planning, route optimization and customer agent interactions - what we might think of as 'last mile, local orchestration.' However, for use cases like end-to-end intelligent omnichannel engagement in high-stakes and heavily-regulated fields like life sciences, relying solely on these capabilities poses significant risks. **Without a centralized orchestration layer that includes a way to capture strategy and goals, and a built-in knowledge base**, autonomous agents could make decisions that are misaligned with strategic goals, lack transparency, and ultimately compromise leadership oversight.

This highlights the critical roles of **Aktana's Optimizer, Orchestrator and Knowledge Nexus.**

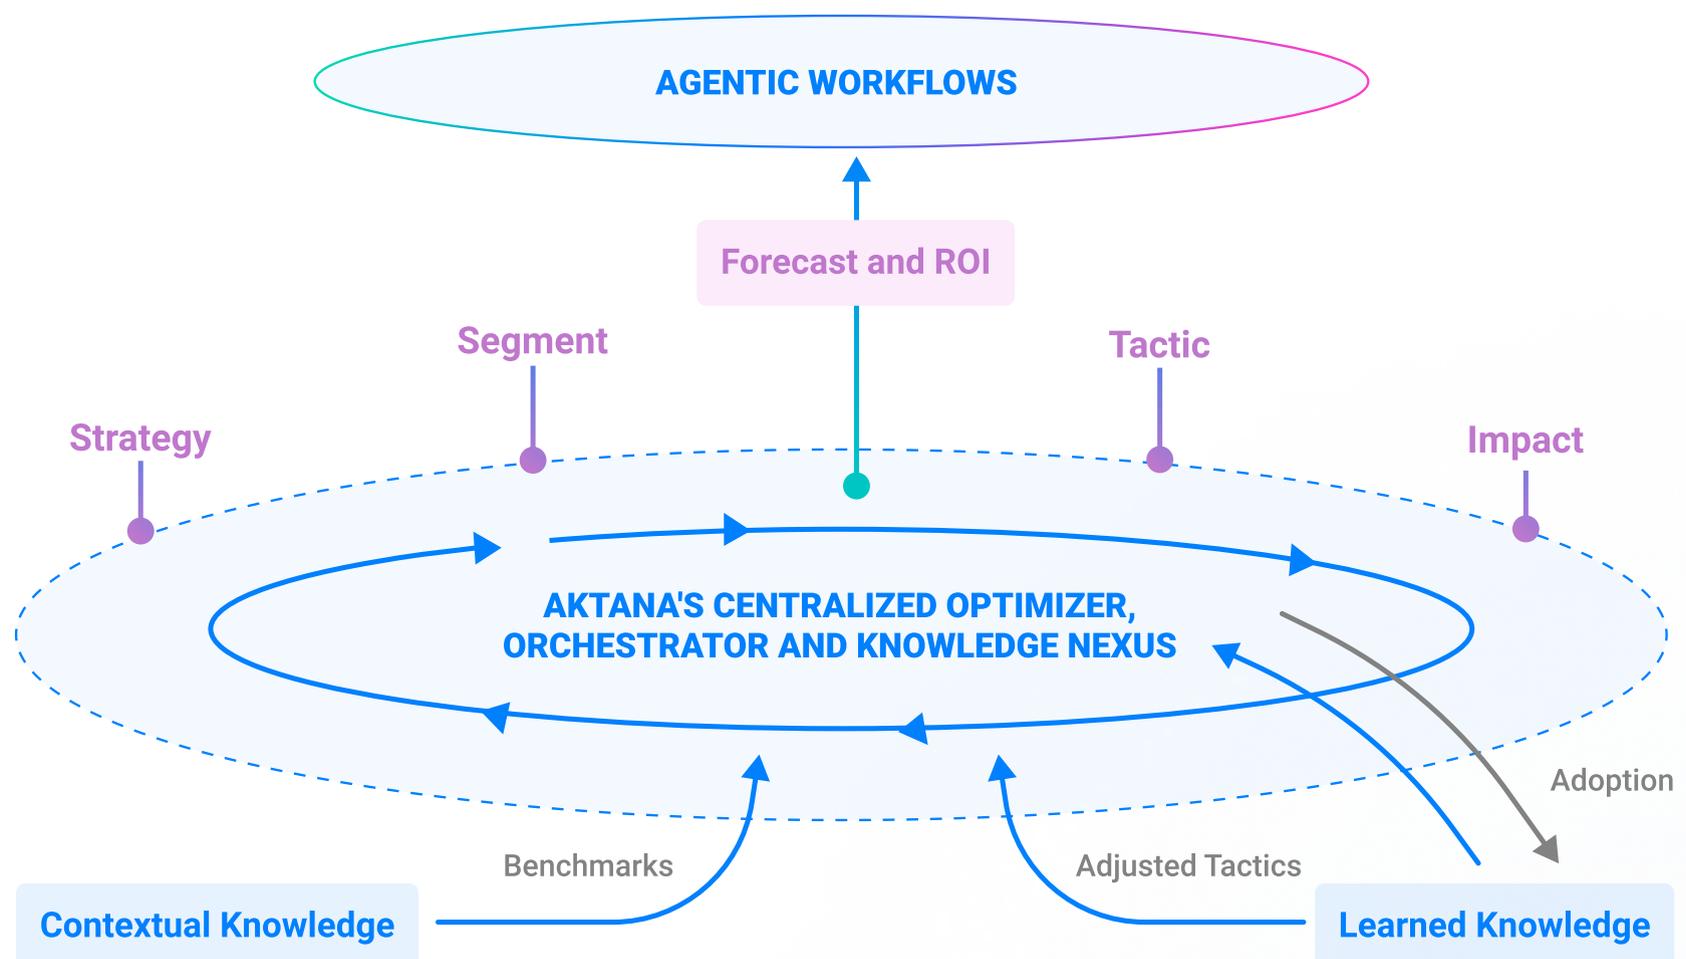
**Aktana's Optimizer & Orchestrator**, function as the centralized intelligence layer that guides Agentic AI, transcending a simple decision engine. Driven by a Dynamically-Constrained Optimization Engine, Aktana calibrates every action in real-time across channels, constraints, and competing priorities and ensures that ROI is always maximized.

The result is a **Scored Forecast**, assigning economic value to each potential omnichannel action. This guarantees that agentic recommendations are proactive but also thoroughly informed, ultimately driving higher-impact business outcomes rather than perpetuating ineffective past strategies.

**Aktana's Knowledge Nexus**, is crucial for Agentic AI, as it provides the necessary context from over a decade of life sciences best practice and experience. With **5,000+ proven tactics and 85+ validated use cases**, it grounds every AI-driven action in real-world success – providing ready-to-use blueprints for boosting tactic effectiveness, accelerating impact, and maximizing ROI. This knowledge base ensures AI agents operate within a structured framework, learn from past strategies, and optimize performance using benchmark data.

### Why This Matters For Life Sciences

Trust is fundamental in life sciences. To ensure that AI decisions are transparent, compliant, and strategically robust, Agentic AI requires integration with a sophisticated intelligence layer. Solutions like Aktana's Optimizer, Orchestrator, and Knowledge Nexus are vital in this regard. We believe that only through such a synergistic combination - where autonomous execution is guided by learned intelligence, contextual knowledge, and global optimization - can AI truly drive transformative, industry-wide advancements.



## Question 7

# How can Agentic AI streamline operations and team workflows in life sciences?

Imagine a world where slow, siloed processes and endless manual reviews no longer hold your teams back. That's the promise of Agentic AI: replacing fragmented workflows with intelligent, adaptive systems that operate in real time. **This includes the ability to dynamically create and deliver personalized communications that resonate with each healthcare professional.**

- **From Manual to Automated Compliance**

Agentic AI goes beyond simply checking for errors - **it actively tailors communications to specific needs and preferences while ensuring continuous compliance.**

A marketing team can now utilize AI to craft diverse, personalized assets that are instantly vetted against business rules, accelerating time to market with communications that are both engaging and fully compliant.

- **From Reactive to Proactive Decision-Making**

Rather than merely responding to lagging indicators, AI can provide near real-time intelligence for all aspects of operations— from clinical trial operations to engagement strategy.

AI could analyze real-time patient data and trial progress to recommend protocol optimizations, thereby shortening trial cycles and accelerating the launch of new brands or indications.

Similarly, if prescribing behavior shifts, field priorities adjust immediately—no need to wait for end-of-quarter analysis.

- **From Siloed to Connected Teams**

Medical, commercial, and regulatory teams often work in parallel, not in sync. Agentic AI connects teams by automatically routing insights and approvals, fostering cross-functional coordination

HCP questions are triaged to the correct expert instantly, improving both response speed and accuracy.

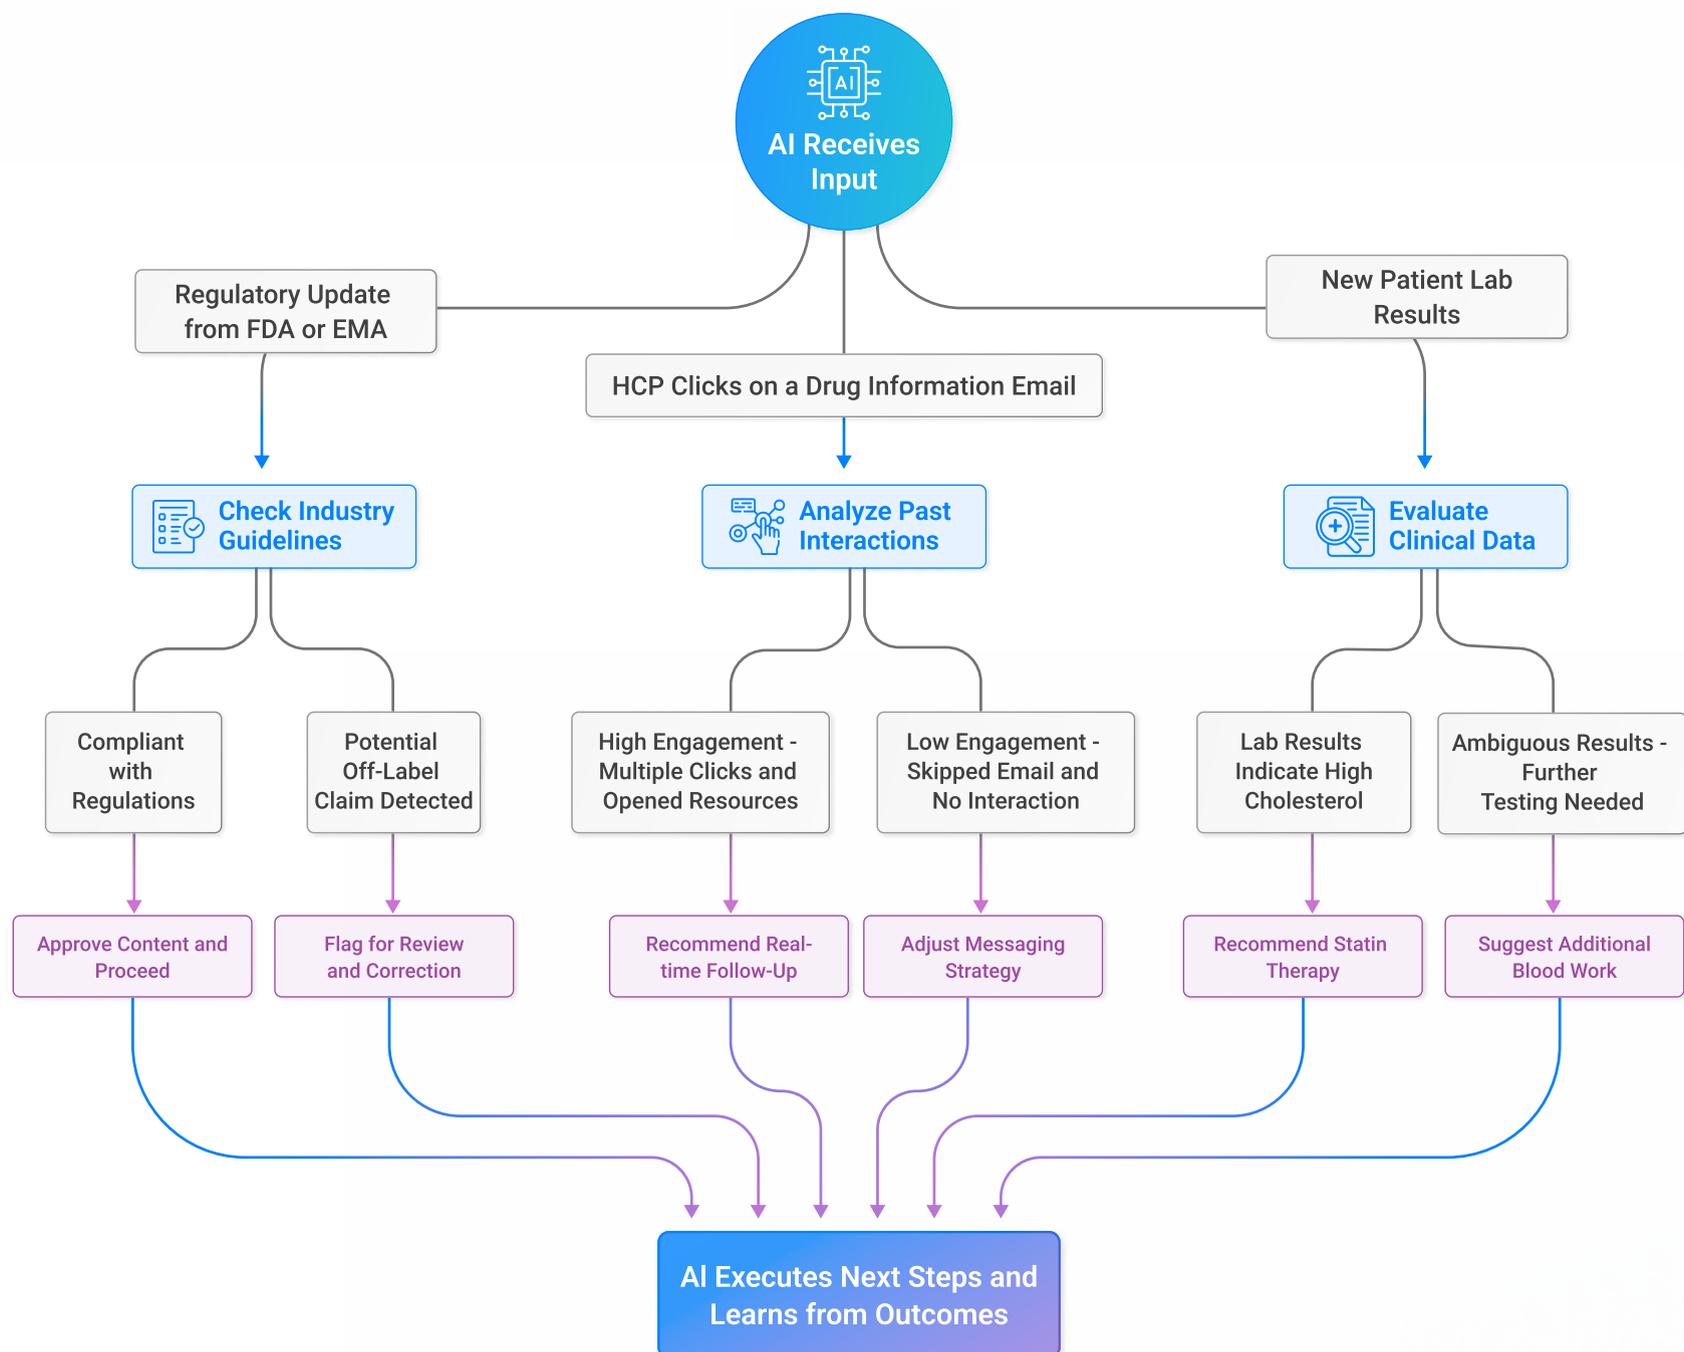
- **From Static to Adaptive Resource Allocation**

Conventional planning is based on historical assumptions. An Agentic AI system can dynamically reallocate resources based on live data.

When competitor activity spikes in a particular region, sales efforts shift accordingly—maximizing impact without added headcount.

Agentic AI transforms automation. By embedding intelligence into every layer of the operation, it can help life sciences shift from reactive execution to strategic coordination, freeing teams to deliver more value.

See the potential in action: The diagram below illustrates how these elements connect. You'll see how Agentic AI processes real-world inputs like regulatory updates, HCP interactions, and patient data to drive decisions and execute actions, all while continuously learning and improving.



**AKTANA**

# **Context-Aware AI for Effective Decision Making**

Agentic AI & Continuous Improvement



# Context-Aware AI for Effective Decision Making: A View from Derek Choy

“ When our team at Aktana expanded our AI systems from pilots to global deployments, we found that scaling these initiatives was more complex than anticipated. Initial attempts relied heavily on manual configurations—rules and static models—which rapidly became cumbersome to maintain at scale. We lacked a **centralized orchestration mechanism** and a structured knowledge base that could dynamically evaluate constraints, and make value driven decisions. This gap led to inefficiencies and stalled adoption.

These experiences taught us the critical importance of having a centralized, dynamic knowledge base at the heart of our AI architecture. Today, Aktana’s multi-agent system leverages this knowledge base to **continuously refine its recommendations, rapidly learning what strategies succeed or fail**. By capturing real-world signals, such as engagement metrics and KPI attainment, it enables proactive adjustments rather than reactive corrections.

The Aktana Strategy Console drives these feedback loops by mapping strategic objectives to tangible KPIs, providing visibility into what’s working and why. For example, stakeholders no longer need to piece together insights from disparate monthly reports; they gain instant strategic clarity to questions like:

- “What is the current progress on our priority KPIs?”
- “Are our current omnichannel strategies effectively driving these KPIs, and why or why not?”
- “What adjustments can we make immediately to optimize performance?”

This enables teams to act swiftly, share best practices, and adapt strategies in real-time for continuous improvement.

Combining traditional machine learning and rules-based systems with newer, more dynamic AI techniques such as reinforcement learning, generative models and Agentic AI has become crucial. This balanced approach ensures recommendations are not only insightful and contextually relevant but also scalable and easy to manage.

Context-aware, continuously learning AI is essential for sustained global impact. At Aktana, we’ve embedded these lessons into our core design, empowering life sciences teams not only to execute today’s strategies more effectively but to dynamically evolve them as conditions change.”

## Derek Choy

Co-founder and Chief Product  
Officer, Aktana

Frequently Asked Questions About

# Agentic AI & Continuous Improvement

Click on a question to jump to answer

**Q1** Does Agentic AI use LLMs and NLP? What role do these technologies play?

---

**Q2** Can Agentic AI solve optimization problems, such as resource allocation in pharma sales and marketing?

---

**Q3** What are some real-world applications of Agentic AI in life sciences, particularly in pharma and biotech?

---

**Q4** What makes a multi-agent system better suited for life sciences than a single-agent AI model?

---

**Q5** What makes context-awareness the key differentiator in AI decision-making?

---

**Q6** Commercial, marketing, and medical affairs teams often work in silos. How can AI-driven agents help them collaborate more effectively?

---

**Q7** How does context-aware orchestration improve HCP engagement?

---

**Q8** Everyone talks about AI improving sales force effectiveness, but what does that actually look like in practice?

## Question 1

# Does Agentic AI use LLMs and NLP? What role do these technologies play?

Yes, **Agentic AI** heavily uses **Large Language Models (LLMs)** and **Natural Language Processing (NLP)** for intelligent decision-making, especially in complex fields like life sciences.

Think of it this way: An Agentic AI system, like a human brain, needs to understand language, reason, and learn from experience.

- **LLMs and NLP (The Brain's Cognitive and Language Centers)**

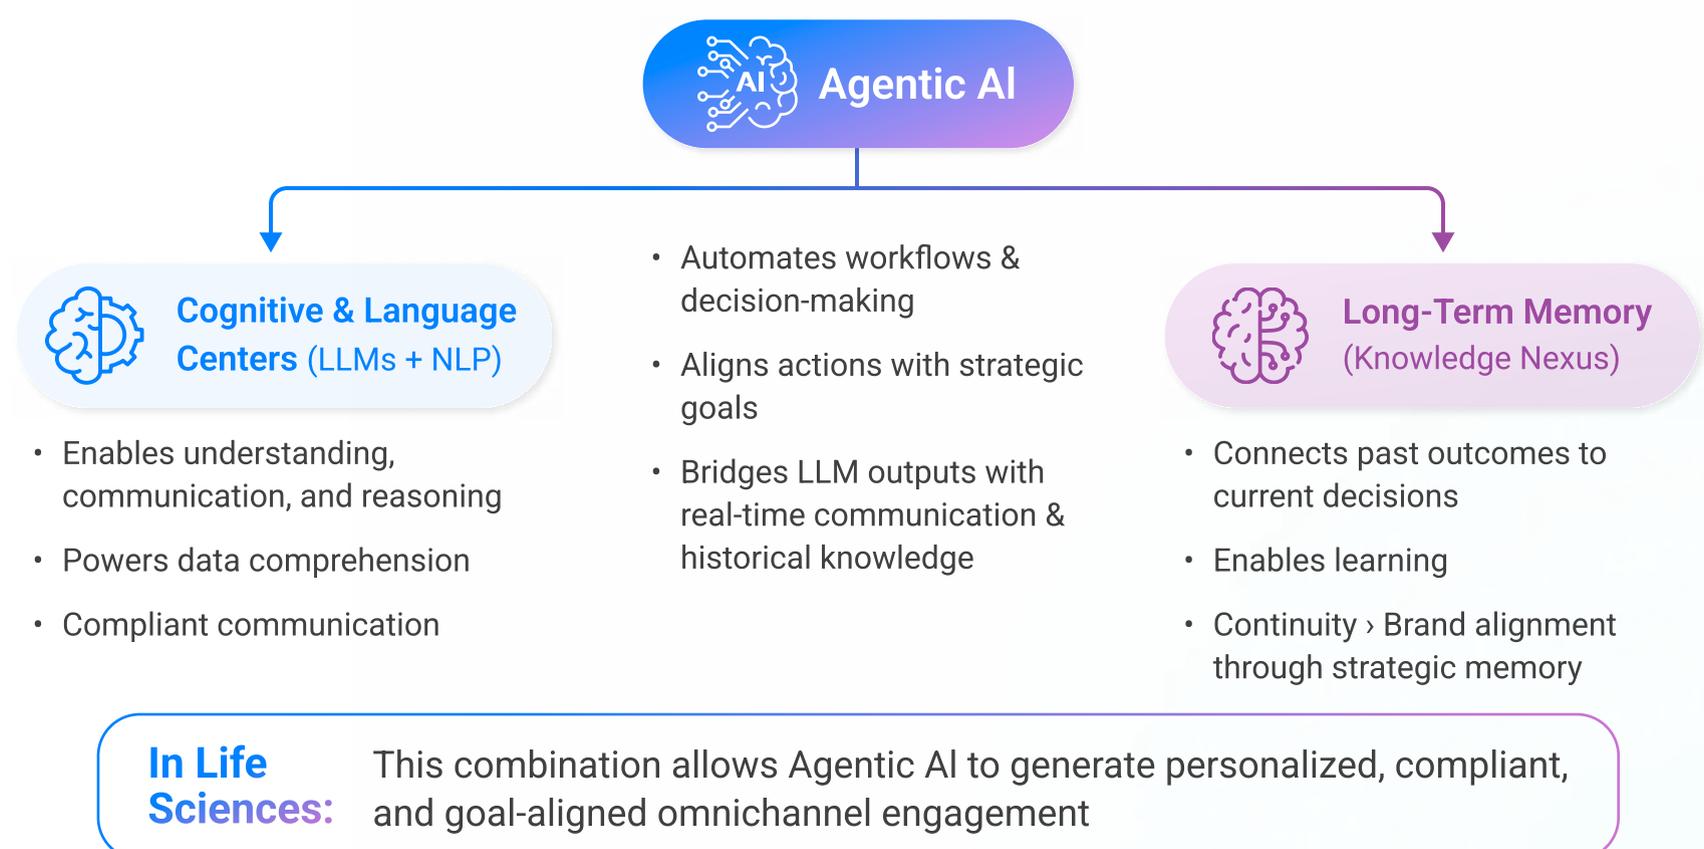
Together, LLMs and NLP form the core “brain” of Agentic AI. They enable the AI to process vast amounts of data, understand context, and generate recommendations. This allows for tailored communication, adjusting content and channels in real-time based on behavior and goals. NLP also helps the AI understand human input and communicate clearly, ensuring messages are compliant and relevant.

However, for true “agentic” behavior, these systems need long-term memory. This is where Aktana’s **Knowledge Nexus** comes in.

- **The Knowledge Nexus (The Brain's Long-Term Memory)**

The Knowledge Nexus acts as the Agentic AI’s long-term memory. It stores successful strategies and lessons learned, connecting current recommendations to past outcomes. This allows the AI to apply accumulated wisdom, ensuring brand strategies can be activated at a quicker pace and performance continuously improves.

In short, LLMs and NLP provide the cognitive power for understanding and communication, while the Knowledge Nexus adds the crucial element of memory and strategic learning. Together, they enable Agentic AI to deliver intelligent, personalized, and goal-aligned orchestration.



## Question 2

# Can Agentic AI solve optimization problems, such as resource allocation in pharma sales and marketing?

Yes. AI-powered agents, operating within an Agentic AI architecture, offer a significant advantage in optimizing resource allocation for pharmaceutical sales and marketing teams. By moving beyond static, rule-based approaches, they enable dynamic, data-driven decision-making that aligns commercial strategies with evolving market conditions and healthcare professional (HCP) preferences.

Here's how these agents address key optimization challenges:

- **Precision in Field Force Deployment:** AI agents analyze prescribing patterns, engagement trends, and competitive intelligence to guide the strategic deployment of sales teams—ensuring the right rep reaches the right HCP at the right time.
- **Adaptive Omnichannel Orchestration:** Agents facilitate the seamless integration and adjustment of digital and in-person interactions, dynamically sequencing channels to meet both brand priorities and HCP behaviors.
- **Data-Driven Marketing and Budget Allocation:** AI systems refine campaign strategies by identifying the most effective messaging, channels, and investment areas—while remaining compliant with regulatory and medical guidelines.
- **Continuous Learning and Agility:** Agents learn from real-time engagement data and user feedback, adjusting their recommendations as market dynamics and brand strategies evolve.

Through these capabilities, an agentic approach can empower pharmaceutical organizations to achieve continuous optimization at scale—driving improved efficiency, stronger HCP engagement, and greater commercial success. Plus, because every recommendation should be explainable and traceable, brand teams can retain full oversight and confidence in how resources are being allocated.

## Question 3

# What are some real-world applications of Agentic AI in life sciences, particularly in pharma and biotech?

Agentic AI is transforming how pharmaceutical and biotech organizations operate—automating compliance, accelerating R&D, optimizing engagement, and improving patient outcomes. Its modular architecture allows specialized agents to collaborate in real time, aligning decisions with strategy, regulation, and human oversight.

Here are some high-impact applications across the life sciences landscape:

- **Regulatory Compliance & Monitoring**

AI agents automatically flag potential compliance risks, track global regulatory updates, and support audit readiness.

**Example:** AI systems review promotional content pre-distribution to detect off-label language or outdated claims.

- **Personalized Medicine & Treatment Optimization**

AI analyzes clinical history, biomarker data, and prior outcomes to support precision treatment recommendations.

**Example:** In oncology, AI-driven platforms assist providers in selecting optimal therapies based on genetic profiles.

- **Drug Discovery & Development Acceleration**

Agentic AI accelerates R&D by identifying promising targets, optimizing molecular structures, and simulating trial outcomes.

**Example:** AI models help reduce preclinical screening timelines by rapidly narrowing down viable compound candidates.

- **Omnichannel HCP Engagement**

Agents orchestrate tailored digital and in-person interactions, aligning content delivery with HCP behavior and preferences.

**Example:** AI curates personalized digital experiences for HCPs, ensuring relevant, timely, and compliant messaging.

- **Sales & Marketing Intelligence**

AI enhances forecasting, territory planning, and rep targeting through real-time market signals and HCP data.

**Example:** AI-powered tools help sales teams prioritize high-value HCPs based on prescribing behavior and digital engagement patterns.

- **Real-World Evidence & Post-Market Surveillance**

AI continuously monitors real-world data to track outcomes, surface safety signals, and inform future strategy

**Example:** AI scans EHRs, claims data, and patient forums to detect emerging trends in adverse events for newly launched therapies.

Across these examples, Agentic AI enables dynamic, explainable, and context-aware decision-making - allowing life sciences teams to move faster without losing control.

## Question 4

# What makes a multi-agent system better suited for life sciences than a single-agent AI model?

In life sciences, decisions aren't linear—they're layered. Choosing the next-best action for an HCP requires balancing brand objectives, engagement history, regulatory boundaries, timing, field capacity, and more. These interdependent variables make the decision space highly contextual—and difficult to navigate with a single, monolithic AI agent.

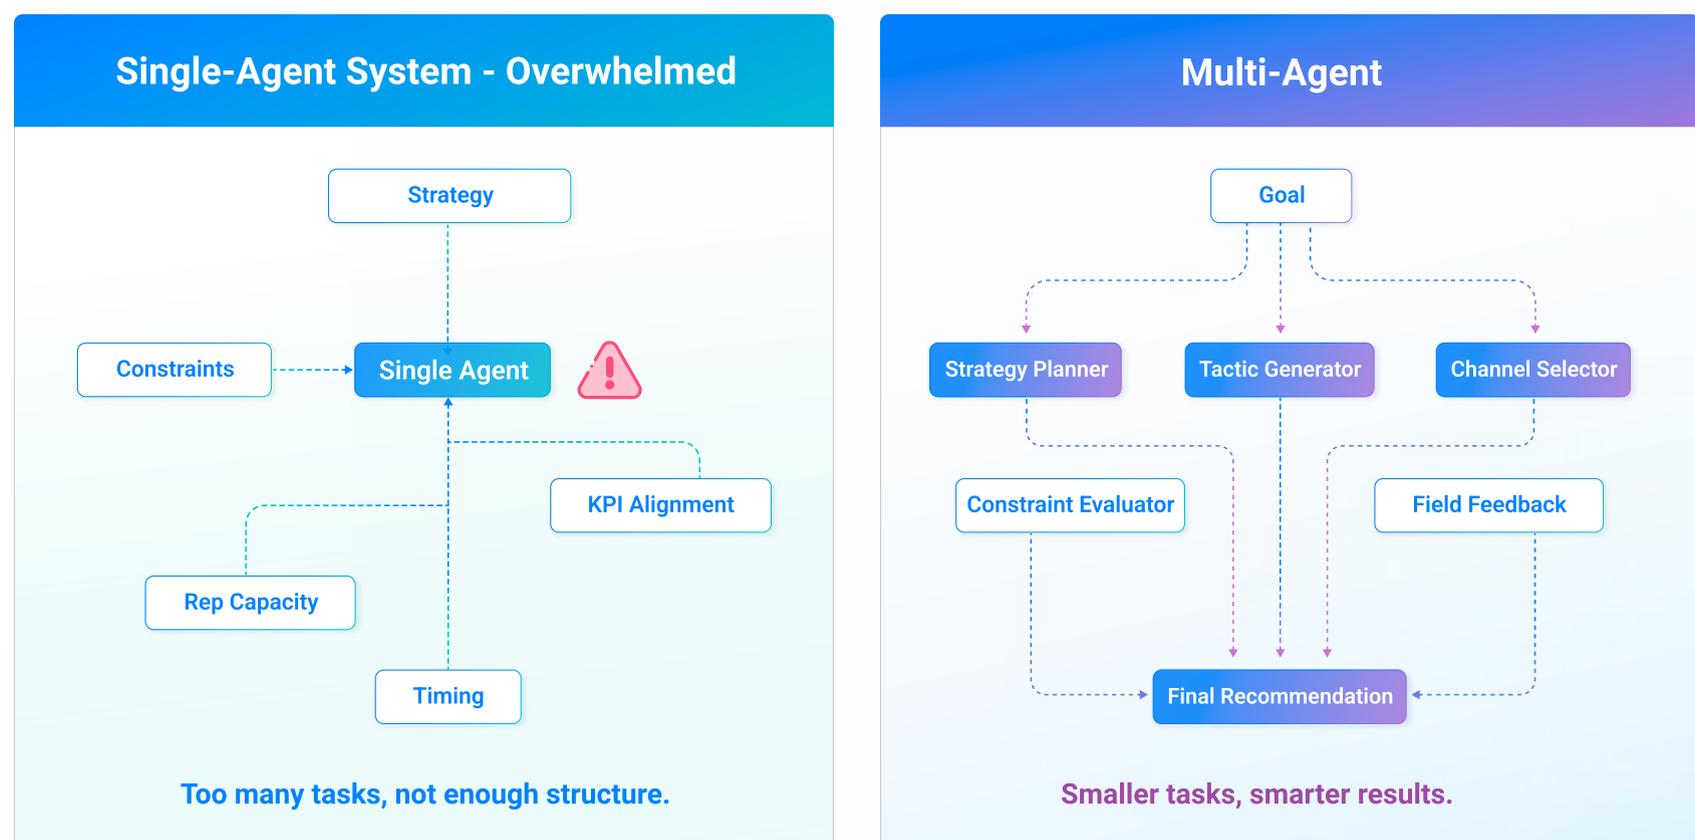
When one AI model is tasked with solving a complex goal like “determine the best action for this provider,” it lacks the structure to reason through competing priorities. The result is often oversimplified outputs or opaque logic, especially when constraints stack up.

That's why the most effective approach is a system of **specialized, collaborative agents**, each focused on a specific aspect of the decision process. One agent might interpret strategic objectives. Another simulates likely outcomes. Others assess channel constraints, timing, or prior engagement effectiveness. This division of labor—akin to **Chain of Thought Reasoning**—lets AI break down complex challenges into smaller, solvable parts while maintaining transparency, intent, and compliance.

Aktana's multi-agent architecture puts this into practice. By distributing intelligence across modular components—each aligned with a business function—our system ensures that recommendations are **explainable, strategic, and real-time**. Every decision reflects not just data, but brand goals, regulatory standards, and field realities.

In high-stakes, regulated environments like life sciences, this isn't just a design preference. It's the only practical path to scalable, trusted AI adoption.

## Why Multi-Agent AI Works (and Single-Agent Doesn't)



## Question 5

# What makes context-awareness the key differentiator in AI decision-making?

In life sciences, static AI models can't keep up with evolving HCP behavior, shifting brand priorities, or real-time market dynamics. **Context-awareness** is what sets Agentic AI apart - enabling the system to adapt on the fly and deliver recommendations that reflect not just strategy, but the reality of the moment.

At Aktana, context-aware decision-making is powered by:

- **The Contextual Intelligence Engine (CIE)**

Our real-time engine continuously synthesizes live data—such as HCP preferences, field activity, and channel availability—to prioritize and deliver next-best actions across both field and digital channels. CIE ensures that every decision is immediately relevant, while still aligned with brand goals. (Read more: [Contextual Intelligence Engine](#))

- **The Knowledge Nexus**

Acting as the system's long-term memory, the Knowledge Nexus captures what's worked (and what hasn't) across brands, markets, and use cases. Tapping into 12+ years of contextual data to inform our agentic decisions, it provides ready-to-use blueprints to boost strategic effectiveness and continuous learning across deployments. (Read more: [Agentic AI](#))

Together, these components ensure that every AI-driven recommendation is timely, explainable, and aligned with brand strategy—even as that strategy evolves. **Because decisions are made transparently and in real time, teams can act with confidence knowing they retain oversight and influence at every step.**

## Question 6

# Commercial, marketing, and medical affairs teams often work in silos. How can AI-driven agents help them collaborate more effectively?

AI-driven agents bridge cross-functional silos by sharing context, coordinating actions, and adapting in real time. They operate with awareness of shared goals, incorporate feedback from all functions, and adjust based on what's actually happening in the field.

By maintaining a **dynamic memory** of past actions, content usage, and planned engagements, agents ensure teams stay aligned even if they operate within different systems or workflows.

For example, if a medical inquiry signals HCP interest in a new indication, an agent can capture that signal and surface it across functions—prompting marketing to update its content arc and commercial to prioritize follow-up. Each function remains informed and responsive, without relying on manual updates or coordination meetings.

Aktana's agentic architecture supports this level of synchronization through its **Optimizer** and Knowledge Nexus - two core components that ensure decisions are both coordinated and contextually intelligent. The Optimizer functions as a dynamic orchestration engine, continuously aligning actions across teams based on real-time priorities, constraints, and strategic goals. Meanwhile, the Knowledge Nexus acts as a shared memory system, capturing insights and outcomes across brand, field, and medical functions to inform future decisions. Together, they ensure life sciences teams engage HCPs with a unified, coherent voice—not three separate ones—while preserving the transparency and control each function needs to stay accountable.

## Question 7

# How does context-aware orchestration improve HCP engagement?

Omnichannel success depends on coordination. That's where AI-driven orchestration delivers the greatest impact. Instead of relying on disconnected campaigns or static plans, AI agents bring structure, timing, and personalization to every interaction.

Agents prioritize actions across field and digital channels based on real-time context: what an HCP has already seen, what they've engaged with, and what's most likely to resonate next. This ensures each touchpoint feels connected and relevant—not repetitive or out of sync.

For example, if an HCP opens a clinical update email, an agent can suppress similar content, queue up a rep visit with complementary materials, or escalate to a higher-value tactic. Every step is sequenced using a full view of the HCP—not just one team's inputs.

At the core of this orchestration is real-time learning in Agentic AI—a feedback-rich approach where every action informs the next. Aktana's multi-agent architecture powers this orchestration through dynamic segmentation, intelligent timing, and continuous learning. And because every action is transparent and traceable, teams can confidently scale engagement—delivering the right message, at the right time, in the right way.

## Question 8

# Everyone talks about AI improving sales force effectiveness, but what does that actually look like in practice?

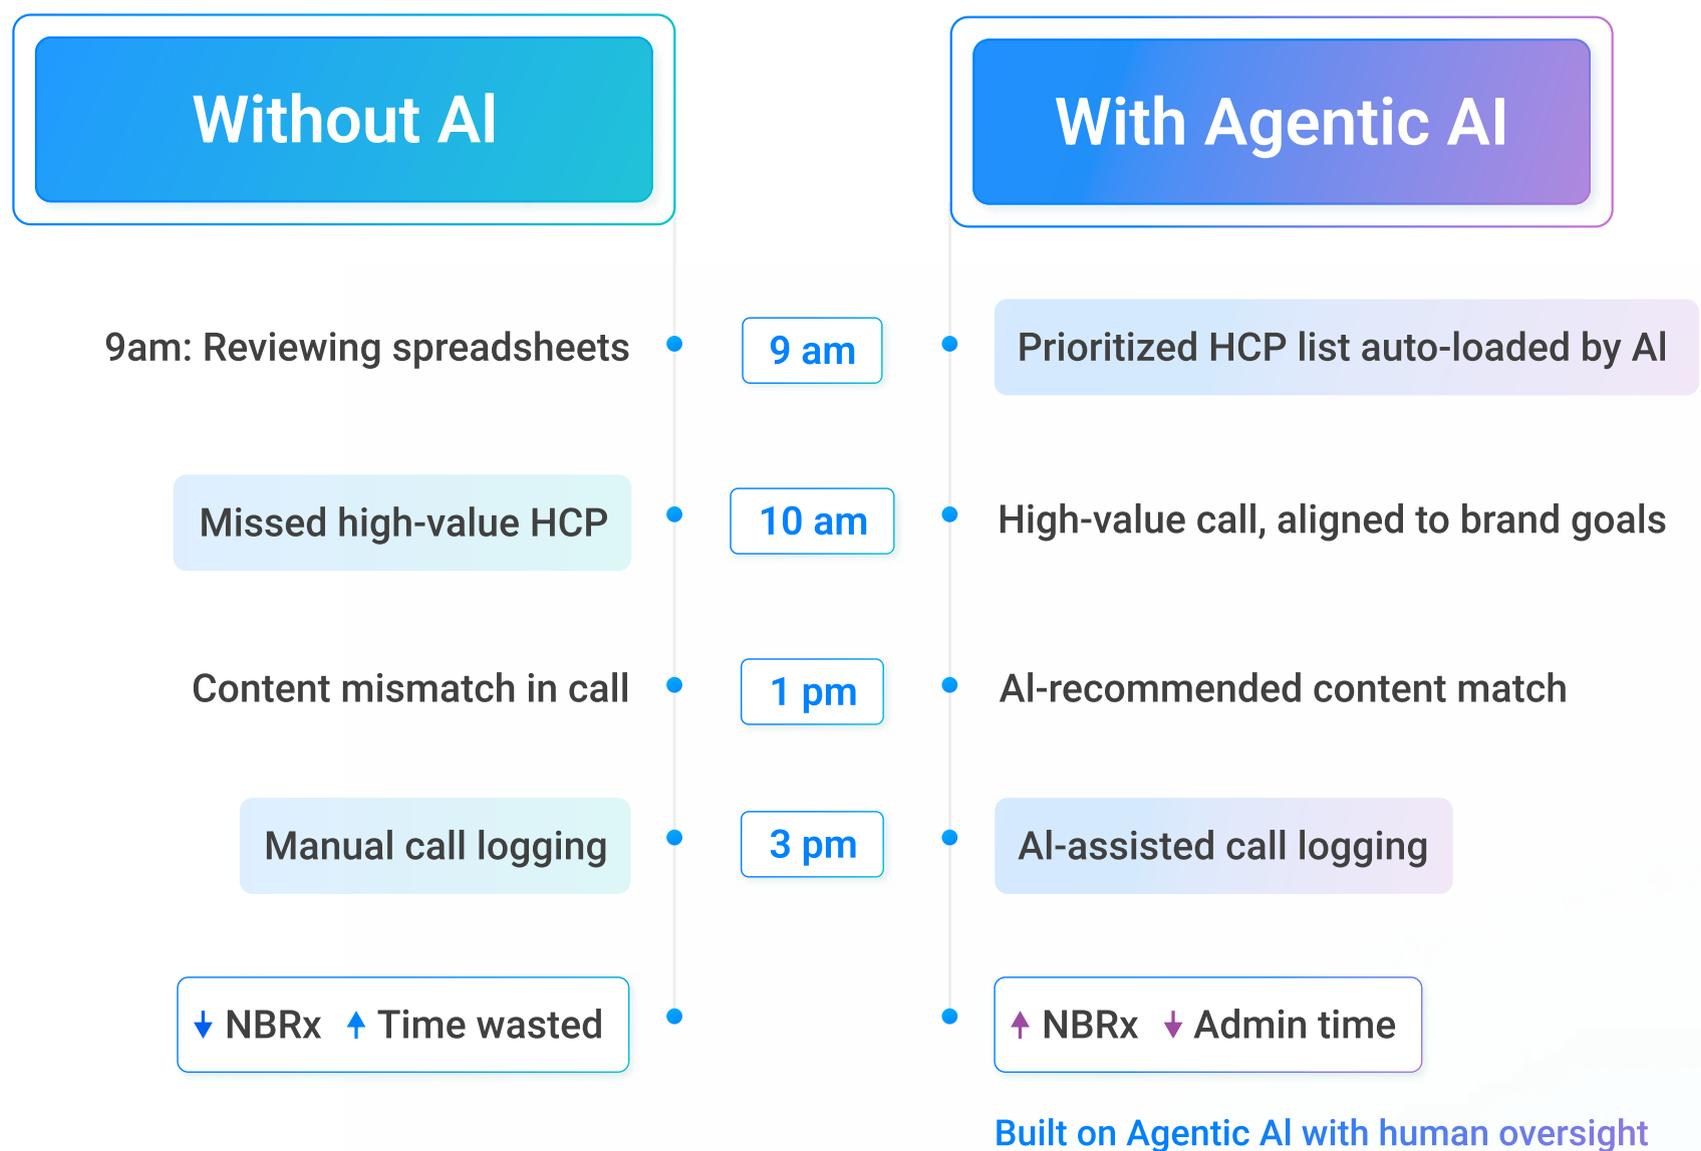
AI improves sales force effectiveness not by adding more tasks, but by helping reps focus on the right ones. From pre-call planning to territory coverage, AI agents streamline decision-making so field teams can spend more time engaging and less time guessing.

Instead of relying on static call plans or gut instinct, reps receive context-aware suggestions: who to prioritize, what to say, which channel to use, and when to engage. These recommendations reflect brand strategy, HCP behavior, rep availability, and compliance rules, ensuring every action is purposeful and aligned.

For example, if a rep has limited capacity this week, AI might prioritize HCPs showing rising engagement trends, suppress outreach to recently contacted providers, and surface pre-approved content tied to brand KPIs—all without requiring manual coordination.

With Aktana, these suggestions aren't just personalized—they're explainable, auditable, and continuously improving. Reps gain confidence, managers see measurable impact, and the entire commercial operation becomes more strategic, agile, and aligned.

## A Rep's Day: Before vs After Agentic AI





# Practicality and Integration as Cornerstones of Future-Ready AI

The pharmaceutical industry's omnichannel landscape is inherently complex, characterized by transitioning CRM systems, siloed field and digital operations, intricate product portfolios, and stringent regulatory constraints. Recognizing these challenges, our approach emphasizes practicality, ensuring that AI adoption integrates seamlessly with existing workflows, systems, and regulatory frameworks.

**Transitioning between CRM systems** can be a daunting task for pharmaceutical companies, often leading to disruptions in operations and an inability to innovate during lengthy migration periods. Aktana's Omnichannel Data Interface (ODI) is designed to support clients through these transitions, facilitating seamless data migration and system interoperability. By harmonizing data across disparate sources, ODI ensures that critical data and intelligence investments remain accessible and actionable during transition, enabling innovation today.

In an industry where organizational change is frequent, maintaining continuity in omnichannel strategies is crucial. Our solutions are designed to adapt to such shifts, ensuring that transitions—whether due to personnel changes, mergers, or restructuring—do not disrupt ongoing operations. By providing a unified, intelligent orchestration platform, we help organizations maintain strategic focus and operational efficiency, even in moments of flux.

The journey toward autonomous AI is not uniform across organizations; it varies based on readiness, infrastructure, and strategic goals. Aktana's agentic AI approach is tailored to meet customers wherever they are on this journey, supporting a full spectrum—from manual campaigns and rule-based execution to highly sophisticated, AI-driven engagement. This phased approach allows teams to gradually build trust in AI technologies while ensuring that each step aligns with real-world capabilities and evolving business needs.

By balancing innovation with integration, and ambition with practicality, Aktana's solutions bridge the gap between today's operational challenges and tomorrow's intelligent automation—paving the way for a future where AI is a trusted, transparent, and indispensable partner to life sciences teams.

Aktana's commitment is to help you navigate and thrive in this evolving landscape practically, transparently, and collaboratively. We believe Agentic AI is more than technology—it's a partner in your strategic vision, helping you achieve ambitious outcomes reliably and compliantly.

We invite you to join us in this conversation, explore further with our comprehensive intelligence briefing, and collaborate with us to shape a future where AI empowers, rather than replaces, human judgment and leadership.

With sincere appreciation and anticipation,



**John Vitalie**  
CEO, Aktana



**Derek Choy**  
Co-founder and Chief Product  
Officer, Aktana

# Ready to learn more about Agentic AI?

Contact our team to get started.

[www.aktana.com](http://www.aktana.com) 1.888.707.3125 [sales@aktana.com](mailto:sales@aktana.com)

**AKTANA**