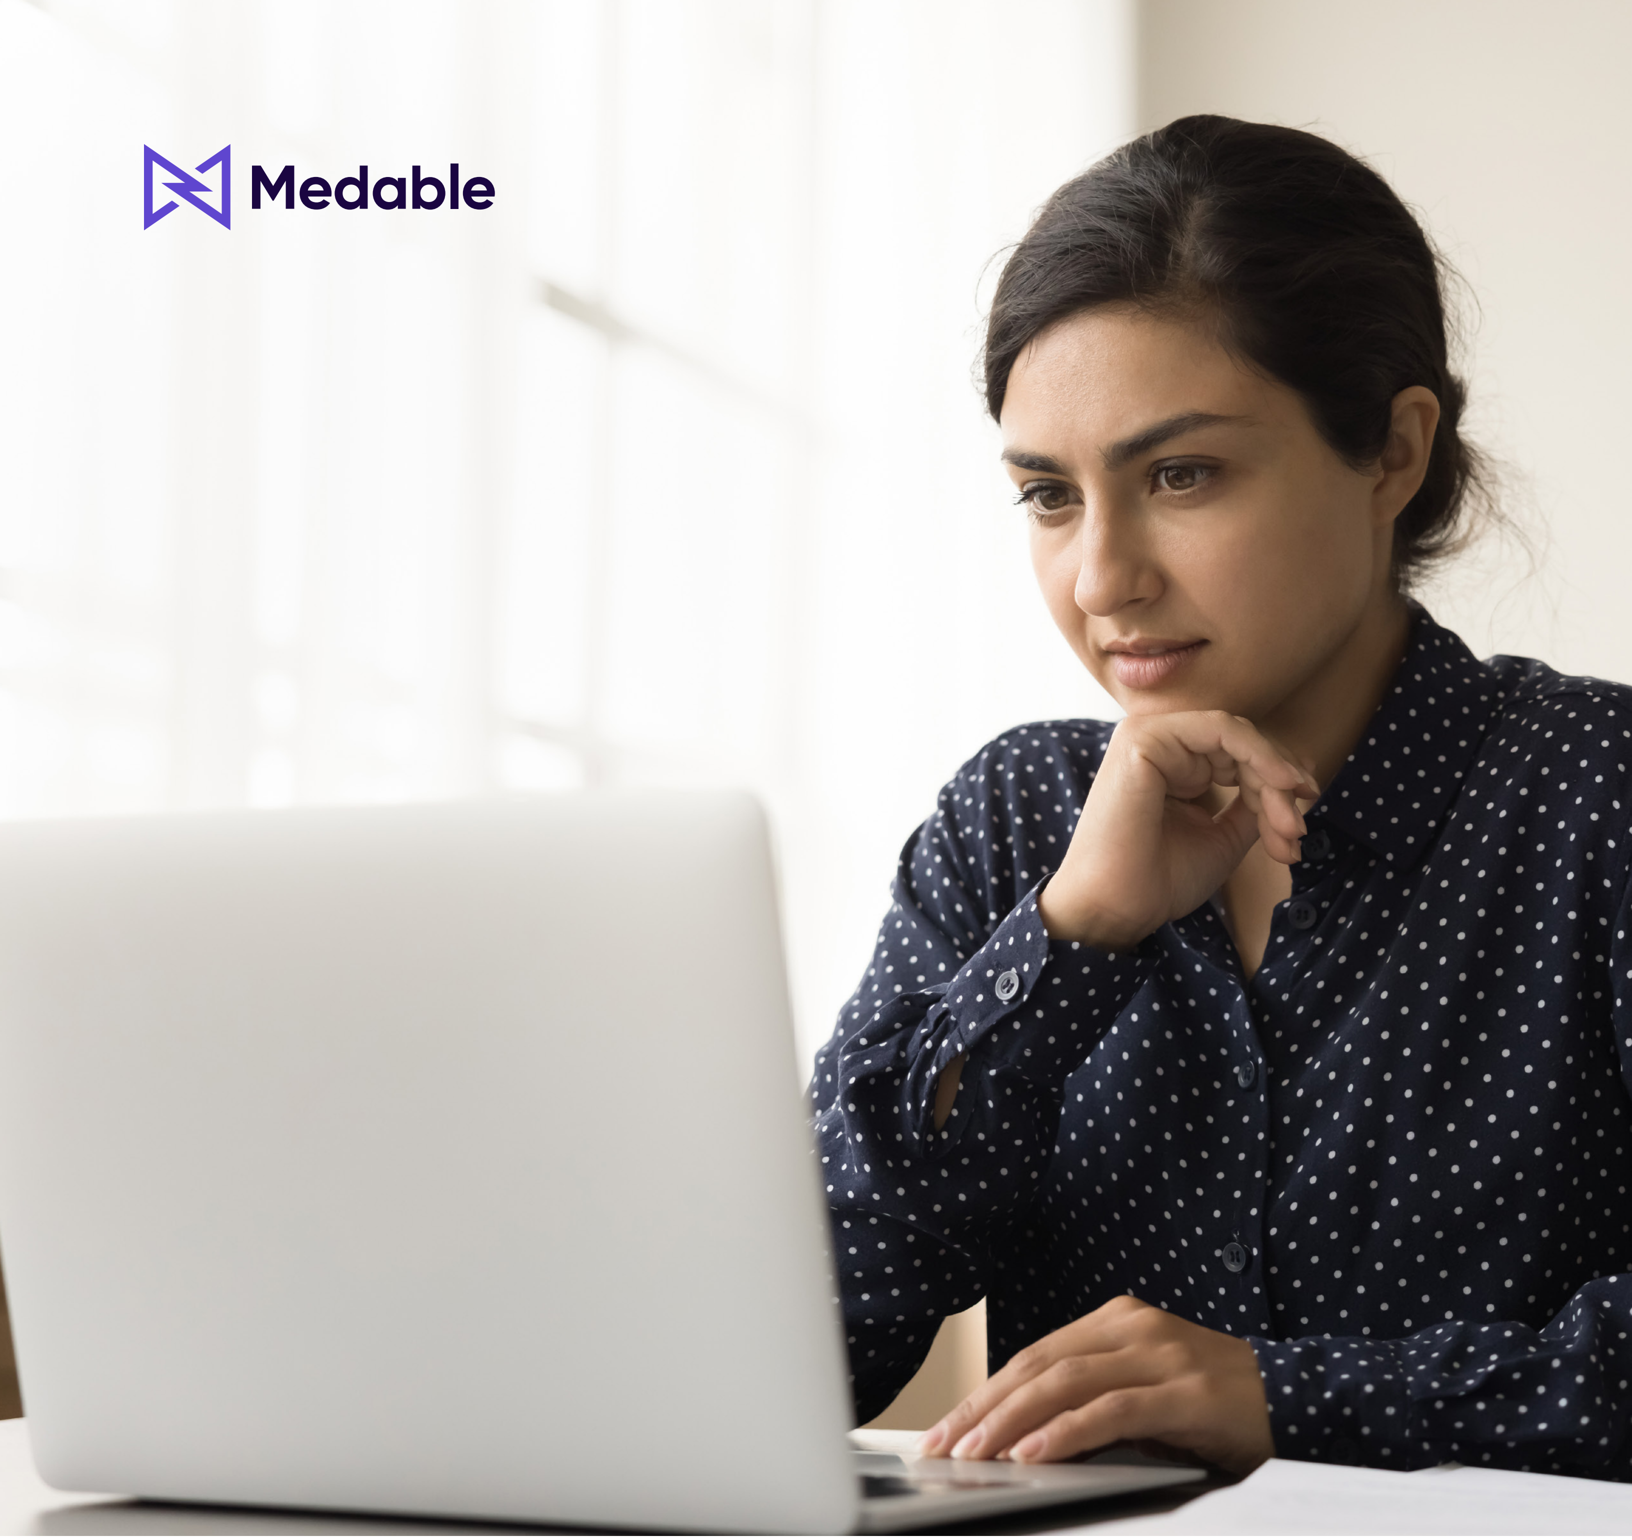

Whitepaper

Eliminate clinical trial white space with the right AI strategy

Table of contents

Executive summary

Clinical trials have become increasingly complex, leading to longer drug development timelines and significant delays for patients. This is due to "white space," which we define here as unproductive time caused by manual, sequential processes and fragmented data systems. A solution lies in agentic AI, a type of AI that can autonomously perform a series of tasks. Agentic AI can reduce white space by:

- Automating administrative tasks and coordinating parallel workflows.
- Integrating data across disparate systems for faster decision-making.
- Monitoring data in real time, allowing human experts to focus on strategic tasks.

However, a human-in-the-loop approach is crucial. Human oversight must be maintained for critical tasks requiring clinical judgment, ensuring regulatory compliance and patient safety. By adopting a specialized AI framework, the industry can transform these delays into accelerated progress, ultimately bringing life-saving therapies to patients faster.

1. Introduction: Reaching the limits of human-only clinical development

It has become clear that our industry has reached the limits of human-only clinical development. As clinical trials have become increasingly complex, the endeavors that people alone can perform are no longer sufficient to generate the momentum needed to address the growing burden of human disease. However, pharma tech is entering a new era defined by rapid innovation, digital transformation, and a re-imagining of what's possible in drug development. The age of artificial intelligence (AI) is here, ushering in a period of rapid evolution as organizations everywhere look to do more with less and respond to a convergence of challenges.

In 2023, Tufts Research quantified the growing increase in clinical trial complexity, stating that the results of their analysis show a "continuing upward trend across all protocol design variables." This is especially true in phase II and III trials, which "average more endpoints, eligibility criteria, protocol pages; investigative sites; countries and data points collected" than other phases.

Additionally, research from IQVIA in 2024 showed that, "on average, new drugs spend 45% of their development time on white space — which IQVIA defines as the time between trial completion and starting the next phase — on the way to regulatory submission."

Unfortunately, a second, lesser known "white space" phenomenon exists as well. This is the unproductive time during individual trial activities, stages, and phases during execution. While this white space has not been studied in detail, the industry knows it exists and is aware of just how much it contributes to the overall duration of each clinical trial.

Much like how clinical trial complexity has grown, clinical trial cycle times also continue to be a challenge. Cycle times are defined as the total duration from the approval of a clinical trial protocol to the database lock (DBL) and the growing challenges in performing each clinical trial and then moving efficiently into the next phase of development are resulting in increasing overall drug development timelines. As reported by Statistica, the average clinical trial cycle from 2020 to 2024 rose by seven months when considering the period from the start of the PhI trial to the conclusion of the PhIII trial.

Average length of a clinical trial cycle from 2020 to 2024

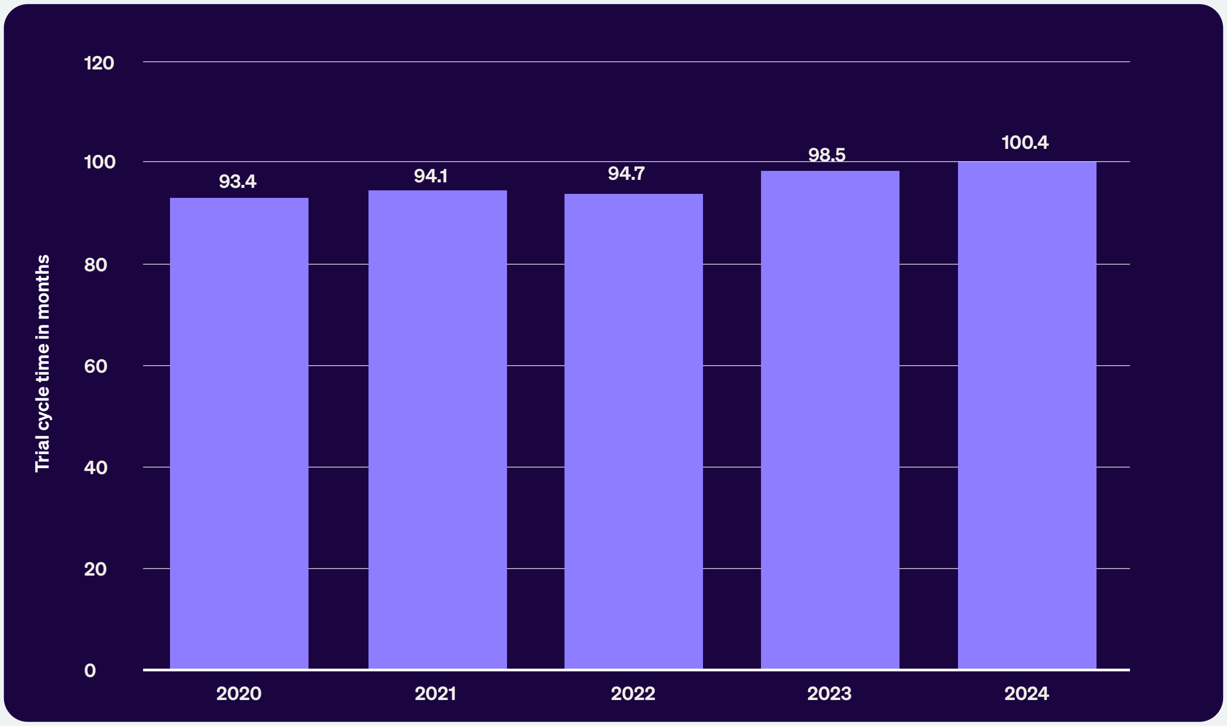

Mikulic, Matej. "Clinical Trial Cycle Length in Months 2024." Statista, 2 May 2025, www.statista.com/statistics/1419920/biopharma-clinical-trial-cycle-time/

2. Examples of white space within trial execution

The unproductive white space that exists during individual trial activities, stages, and phases has several examples that are well-known within the industry.

Between protocol finalization and site activation

- Waiting on regulatory, EC/IRB and site committee approvals
- Incomplete or unclear site feasibility feedback
- Misalignment between operational readiness and document readiness
- Last-minute protocol amendments impacting eClinical system build
- Delays in site training or IRB approvals

Between first activation and last patient last visit (LPLV)

- Site activation ramp up as each country and site comes online
- Delays and lulls in patient enrollment as sites struggle to identify and bring in participants
- Protocol amendments resulting in stagnation, recruitment pauses and retraining
- Decision making delays due to lack of data cohesiveness across clinical systems at critical junctures such IDMC reviews, cohort transitions and other critical decision time points

Between last patient last visit (LPLV) and database lock

- Data cleaning bottlenecks and query resolution after final patient visits
- SAE reconciliation and endpoint adjudication
- Waiting on central lab, imaging reads or other external data sources

These pervasive bottlenecks throughout the clinical trial lifecycle (from protocol finalization through database lock) stem from three fundamental systemic issues that plague traditional clinical development operations.

Manual processes create dependency on human availability and hand offs, whether it's manually tracking regulatory submission status, coordinating site feasibility assessments, or resolving data queries one by one after patient visits conclude.

Siloed systems exacerbate these delays by fragmenting critical information across disconnected platforms: regulatory teams work in isolation from operational teams, site feasibility data doesn't integrate with document management systems, and post-study data cleaning occurs separately from ongoing safety monitoring, preventing real-time visibility, and coordination.

Sequential workflows force activities that could occur in parallel to wait in lengthy queues. For instance, protocol amendments trigger cascading delays through eClinical system builds, regulatory approvals must complete before site training can begin, and database lock cannot proceed until every query is resolved and every central lab result is received. This combination creates a compounding effect where delays in one area cascade through dependent processes, turning what should be efficient, parallel operations into lengthy, fragmented sequences that can extend trial timelines by months or even years.

Each of the white space examples above can increase clinical trial cycle times, and have been the focus of many functions within the pharmaceutical, biotech and CRO space. In fact, IQVIA's research notes that advancements in clinical trial process and technology have yielded an average reduction of seven months in the duration of actual clinical trials. However, "for trials completed in the past five years, the white space between these trials increased 14 months, offsetting the seven-month reduction in trial durations and resulting in seven months longer overall durations."

The industry should recognize a key risk: rushing through early trial stages often backfires. Moving too fast can create data and system integration problems that ultimately take more time to fix than the speed was meant to save. This problem often occurs when scaling human-intensive tasks, which is why shifting the scaling burden to technology can improve overall system efficiency.

White space between phases by therapeutic area

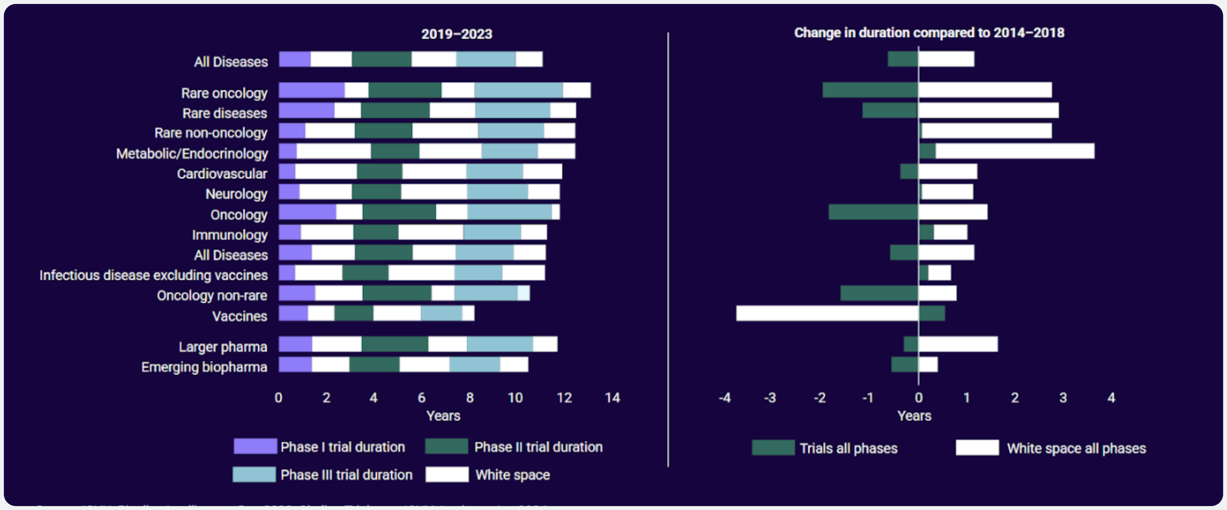

Source: IQVIA Pipeline Intelligence, Dec 2023, Citeline Trailtrove, IQVIA Institute, Jan 2024

The result of these increased cycle times is that patients often wait years for therapies to progress through the clinical pipeline, not because of scientific limitations, but due to entrenched inefficiencies in how trials are planned and executed. Thankfully, industry technology has advanced enough to offer sponsors and CROs new, more effective ways to tackle these issues.

3. The many types of AI and their promise in clinical trials

Before we address how AI can help clinical trials, it's important to define what we're talking about. This is especially true as artificial intelligence is a rapidly evolving field. While the term "AI" is often used as a catch-all, it actually encompasses several distinct types and categories.

AI and machine learning: A parent-child relationship

At the most fundamental level, AI is the broad science of creating intelligent machines. Machine Learning (ML) is a core subfield of AI. It's a method where systems learn from data, identify patterns, and formulate decisions with minimal human intervention.

Generative AI: Creating something new

Generative AI is a class of AI models that can create new content, such as text, images, music, or code. Unlike traditional AI that primarily analyzes and classifies existing data, generative AI takes what it has learned from a vast dataset and produces original, realistic-looking outputs.

Agentic AI: The future of automation

A newer and increasingly important concept is Agentic AI. While a Generative AI model simply responds to a prompt and is therefore reactive in nature, Agentic AI can autonomously perform a series of steps to achieve a larger goal. It can think, plan, and take actions on its own. For example, an AI agent could be tasked with "planning a trip to New York." It would then autonomously break down the task into sub-tasks: searching for flights, booking a hotel, and creating an itinerary, all while adapting to new information in real-time.

Horizontal vs. vertical AI: The scope of an AI company's application

The AI companies that exist today can be categorized by the breadth of their application focus.

Horizontal AI companies develop general-purpose solutions that can be applied across multiple industries and use cases. These companies typically offer foundational technologies—such as large language models like GPT or Claude. These can be utilized by law firms, marketing agencies, healthcare providers, and countless other sectors. The value of horizontal AI companies lies in their versatility and broad market reach, providing scalable solutions that serve diverse customer bases.

Vertical AI companies, in contrast, focus on building solutions specifically tailored for a single industry or highly specialized function. Importantly, these companies often utilize the same underlying AI technologies (such as GPT-4 or other foundational LLMs) as their horizontal counterparts, but differentiate themselves through specialized grounding, domain-specific context, and carefully crafted frameworks that optimize performance for their target sector. Furthermore, vertical AI companies understand the implications of using technology platforms and systems within their specific industry and the challenges, such as system validation, that are present. The deep grounding of solutions from vertically focused AI companies, and the development of approaches that allow for the necessary validation documentation, provide advantages to the real world use of these models in a highly regulated environment.

For example, a vertical AI company in healthcare might use the same core language model as a horizontal provider, but enhance it with medical terminology, clinical workflows, regulatory knowledge, and healthcare-specific data sets.

This specialized grounding allows the AI to understand the nuances of medical documentation, regulatory requirements like FDA submissions, or clinical trial protocols—context that a general-purpose model would lack.

Because of this deep specialization in grounding and context, vertical AI solutions often significantly outperform horizontal alternatives within their specific domain, despite potentially using the same foundational AI technology. The vertical company's greatest strength lies not in the AI model itself, but in the specialized expertise, data, and frameworks they've built around it to serve their target industry's unique needs.

4. How AI eliminates clinical trial white space

Agentic AI offers a compelling path forward by directly addressing the operational characteristics of clinical development work. Many tasks within a trial are repeatable, rules-driven, and data-intensive, which is ideal for automation and also what makes the tasks all the more challenging to perform at scale with human resources. However, to be successful within such a challenging, and highly regulated environment, these types of activities will require a more nuanced understanding of context and are therefore more adept for a vertical AI.

Agentic AI agents are built to handle exactly this blend of complexity and repetition. They can pull information from multiple systems, apply relevant rules or guidelines, and execute decisions or trigger workflows without requiring human initiation.

By working continuously and without fatigue, agents ensure that no day is lost to waiting for the "next available" person to process a task.

Keeping a human in the loop

Agentic AI systems, while powerful, are not inherently suited to fully autonomous operation in all regulated contexts, particularly as organizations are still learning to effectively collaborate with these technologies. Current limitations in large language models require active human supervision to ensure critical outputs meet the exacting standards of regulatory bodies. This oversight becomes even more crucial when considering that regulatory violations can result in severe consequences including product delays, financial penalties, and potential harm to patients.

Tasks appropriate for AI autonomy

Certain routine, well-defined tasks can be effectively delegated to agentic AI systems with minimal human intervention. This is especially efficient within clinical research where data is often stored across more than 15 disparate systems, requiring significant effort from humans to aggregate and analyze this data. An agent with access directly to this data can perform this analysis in seconds, seamlessly combining data to identify the correct action to take.

Administrative and follow-up activities

- Following up with investigational sites regarding outstanding documents and submissions
- Tracking and querying data entry timelines and milestone completions
- Monitoring protocol deviations and sending automated reminders for corrective actions
- Managing routine correspondence with study coordinators and site personnel
- Generating standard progress reports and dashboard updates
- Scheduling routine meetings and coordinating calendar management across multiple stakeholders

Data processing and analysis

- Performing initial data cleaning and validation checks
- Generating preliminary statistical summaries and trend analyses
- Creating draft patient recruitment and enrollment reports
- Conducting routine database queries for operational metrics
- Preparing standardized interim safety data summaries

5. Tasks requiring critical human oversight

However, several mission-critical activities must maintain robust human oversight to ensure regulatory compliance and maintain the integrity of the research process:

Regulatory submissions and documentation

- **SAE (Serious Adverse Event) reporting:** Human clinicians must review, interpret, and approve all serious adverse event reports, as these directly impact patient safety and require clinical judgment that AI cannot reliably provide
- **Regulatory submission documents:** Creation of CTD (Common Technical Document) sections, INDs (Investigational New Drug applications), and other regulatory filings require human expertise to ensure accuracy, completeness, and compliance with evolving regulatory guidance

- **Protocol amendments:** Any changes to study protocols must be reviewed by qualified medical and regulatory professionals to assess clinical implications and regulatory requirements

Clinical data management

- **EDC query creation and resolution:** While AI can identify potential data discrepancies, creating queries into Electronic Data Capture systems requires clinical understanding to ensure queries are clinically relevant and appropriately prioritized
- **Medical coding and review:** Assignment of adverse event terms and concomitant medication coding requires medical expertise to ensure clinical accuracy
- **Data lock and database release:** Final database reviews and approvals require human sign-off to ensure data integrity and completeness

Quality assurance and compliance

- **Audit preparation and response:** Interactions with regulatory inspectors and preparation of audit materials require human judgment and regulatory expertise
- **Risk assessment:** Clinical and regulatory risk evaluations must incorporate human experience and clinical judgment that extends beyond what AI can currently provide
- **Corrective action plans:** Development of CAPA (Corrective and Preventive Action) plans requires human understanding of systemic issues and organizational capabilities

6. Implementing effective human-AI collaboration

The most effective approach involves establishing clear governance frameworks that define decision boundaries, approval hierarchies, and escalation procedures. Human oversight should be risk-proportionate, with more critical processes requiring senior-level review and approval, while routine tasks can be managed by junior staff with AI assistance. Regular training and competency assessments ensure that human supervisors maintain the skills necessary to effectively evaluate and guide AI-generated outputs.

This human-in-the-loop model not only ensures regulatory compliance but also serves as a critical learning mechanism, allowing organizations to gradually expand AI autonomy as they develop better understanding of system capabilities and limitations within their specific regulatory context.

7. Examples of white space elimination with agentic AI

The application of agentic AI to site activation processes demonstrates the transformative potential of this technology in addressing one of the most persistent sources of white space in clinical development. Traditional site activation involves numerous sequential steps, including site identification and qualification, contract negotiation, regulatory submissions, training delivery, and technology deployment. Each of these steps typically requires coordination between multiple stakeholders, with hand offs between different functional groups creating opportunities for delays and communication breakdowns. Agentic AI systems can coordinate these activities in parallel rather than sequentially, automatically preparing site packages while contract negotiations are proceeding, scheduling training sessions based on real-time availability data, and proactively identifying potential obstacles that might delay site activation.

Agentic systems can perform continuous, real-time data monitoring across multiple data sources, automatically identifying patterns that suggest data quality issues, protocol deviations, or safety concerns. Over recent years CRAs have become bogged down with multiple eClinical systems collected data from sites and patients, reducing their ability to provide guidance and support to the sites running the study and reducing them to chasing missing data and trying to provide oversight across multiple systems. Rather than waiting for scheduled monitoring visits or periodic data reviews, AI Agents can provide continuous oversight that enables proactive intervention before issues become significant problems. The abilities embedded in these systems enables them to learn from previous site activation experiences, identifying patterns that predict which sites are likely to experience delays and what interventions are most effective in addressing specific types of obstacles. Rather than treating each site activation as an independent project, agentic systems can leverage accumulated knowledge to optimize approaches for different types of sites, geographical regions, therapeutic areas, and protocol complexities. This learning capability means that system performance improves continuously as more data becomes available, creating compound benefits that extend beyond individual trials to impact entire development portfolios.

The data management and monitoring capabilities of agentic AI systems offer perhaps the most immediate opportunities to reduce white space in ongoing trials. Traditional data monitoring requires human reviewers to examine individual patient records, often requiring review of data across multiple systems, identify potential issues, and coordinate with sites to resolve queries or discrepancies. This process, while necessary for ensuring data quality, creates significant delays in database lock activities and can become a bottleneck in study completion timelines.

Finally, the integration capabilities of agentic AI systems address one of the fundamental structural barriers to reducing white space; decision making. These systems can operate across traditional organizational boundaries, accessing data from multiple sources, coordinating activities between different functional groups, and providing comprehensive visibility into trial performance that enables more effective decision making.

Rather than requiring human intermediaries to translate information between different systems and stakeholders, agentic systems can serve as intelligent integrators that ensure information flows seamlessly throughout the trial ecosystem. An agent created towards eliminating white space during study startup can quickly synthesize regulatory requirements, therapeutic area precedent, and sponsor objectives to draft protocol designs and trial documents in record time.

Across all these examples, agents replace idle time with active progress, ensuring that the moments between trial milestones are used to advance the program rather than to wait for the next task to begin.

8. A framework to bring the right AI to your organization

Successful AI agent adoption begins with early stakeholder and executive buy-in. Involving executives, functional leaders, and end-users from the outset ensures alignment on vision, priorities, and success metrics. Early engagement helps secure support, amplifies sponsorship across the organization, and ensures that critical insights are heard. Once sponsorship is in place, organizations should define objectives and scope clearly. This includes differentiating between using AI agents as augmentors, assisting human roles, versus leveraging them for more substantial role automation or replacement. It is essential to articulate the scope of agent roles, boundaries, and interactions to all stakeholders, using intuitive workflows and transparent configuration settings to avoid confusion.

Equally important is ensuring regulatory and ethical alignment. By adopting a compliance-first approach, organizations can align with industry standards such as ISO 27001, SOC 2, HIPAA, GDPR, and 21 CFR Part 11. Human oversight must remain central, with clear accountability, well-defined permission settings, and reliable audit trails. With governance in place, the next step is to prepare a gradual implementation plan. Organizations can pilot with low-risk use cases to demonstrate value and refine processes. Flexible deployment tools and Human-in-the-Loop features make incremental adoption easier. By taking an iterative approach, teams can progressively expand agent capabilities based on performance insights and feedback, steadily increasing autonomy at a pace that matches organizational readiness.

Adoption also requires champions who can lead the way. Identifying and empowering internal advocates early in the process helps drive adoption, address concerns, and build confidence across teams. At the same time, organizations must proactively address resistance by openly discussing the role of agents, showcasing positive real-world use cases, and setting realistic expectations. A user-centric design and transparent communication of both opportunities and limitations help establish trust and demonstrate immediate value, making adoption feel natural and beneficial from the start.

Educating stakeholders on human-agent collaboration ensures that AI adoption is framed as augmentation rather than mere automation. By positioning agents as collaborative tools that enhance, not replace, human decision-making, organizations foster trust and engagement. Transparent reasoning capabilities, such as explainability features and clear chains of thought, help ensure dependable collaboration and accountability. Alongside education, training and communication are vital. Comprehensive training programs, supported by user-friendly interfaces and extensive documentation, simplify onboarding and ensure that teams understand agent capabilities, limitations, and operational boundaries.

The world of AI development is vast, but it's crucial to draw a distinction between AI for other industries and AI for specific, hard-to-get-into verticals. Avoid choosing a broad company that does not have the grounding within your specific vertical. A partner with deep expertise in clinical trials, and ideally your specific TA and trial type, is necessary in order to understand the unique challenges you will face.

For example, consider the writings of York.IE, an advisory and investment firm that helps technology companies grow. York.IE states:

> "One of the biggest limitations of horizontal AI is its lack of domain-specific expertise. A generic model like ChatGPT can generate a broad range of responses, but without access to proprietary industry data, it often fails in specialized use cases."

This expertise is vital for navigating complex issues like AI council approvals or industry-specific compliance requirements. A generic AI company may offer a platform, but a specialized partner offers a proven framework that is pre-validated for the nuances of your industry, ensuring long-term success and a more robust solution.

Finally, ongoing monitoring, measurement, and optimization help sustain long-term success. Continuous feedback loops enabled by real-time monitoring and performance analytics allow organizations to refine and improve their approach. Careful tracking of performance metrics and benchmarks provides accurate insights into agent performance and impact over time. Together, these steps create a structured, ethical, and effective pathway for AI agent adoption within organizations.

9. Conclusion: From white space to breakthroughs

The persistent white space in clinical trials is a solvable problem. By leveraging a purpose-built AI framework, organizations can transform idle time into active progress, accelerate decision-making, and unlock capacity that has long been buried under administrative delays.

The right AI strategy to accomplish this is based on your needs and pairing the right problem definition with the right vertical expertise, while deploying solutions that are both operationally viable and scientifically sound. Agentic AI, in particular, holds the promise of delivering continuous momentum across every trial phase, ensuring that the next critical step is never waiting on the last to conclude.

By adopting a framework-driven approach to AI, anchored in the realities of clinical research, leaders can take decisive steps toward eliminating these gaps once and for all.
