

PHUSE CONNECT 2026

Transforming Clinical Trials

Present Innovations and Future Paths
with AI and Agentic Tools

Sandeep Purwar

Summit Therapeutics, Palo Alto, USA

March 2026

OVERVIEW

Agenda

01

The Clinical Trial Challenge

Understanding pressures and complexity driving AI adoption

03

AI Maturity Model

From assistive to agentic: the evolution pathway

05

Tech Industry Examples

Lessons from software development and DevOps for Biotech adoption

07

Implementation Roadmap

24-month phased adoption approach

02

AI Today

Current applications across study design, data management, and submissions

04

Latest Agentic AI Developments

Market trends and enterprise adoption 2024-2025

06

Agentic AI: 3-5 Years

Near-future capabilities and transformation impact

08

FDA-Aligned Framework

Governance, validation, and compliance requirements

The Clinical Trial Challenge

BioPharma faces unprecedented complexity, cost pressures, and data explosion



Increasing Complexity

More global trials, adaptive designs, and diverse endpoints requiring sophisticated coordination across multiple sites and regulatory regions



Rising Cost Pressures

Demanding timelines and budget constraints requiring operational efficiency without compromising quality or compliance



Data Explosion

Volume and variety growing exponentially from EDC, labs, imaging, ePRO, wearables, and real-world evidence sources

? Three Critical Questions for Responsible AI Adoption

1 **Where are we today?**
Understanding current AI capabilities

2 **What's possible in 3-5 years?**
Realistic agentic AI potential

3 **How do we get there safely?**
Compliant implementation path

CURRENT STATE

AI Today: Four Critical Application Areas

Current AI is assistive and co-pilot, augmenting human expertise rather than replacing it



Study Planning & Documentation

- ✓ **Draft protocol content:** Generate first drafts of protocol sections, SAPs, and CSR sections from structured inputs
- ✓ **Harmonize documents:** Propose edits for consistency across protocol, SAP, and regulatory documents
- ✓ **Prepare materials:** Create annotated CRFs and quality checklists for trial documentation

 **Human role:** Scientific correctness & regulatory alignment



Data Management & Cleaning

- ✓ **Suggest queries:** Identify discrepancies and prioritize data queries based on patterns and historical trends
- ✓ **Detect anomalies:** Flag inconsistent patterns across visits, sites, and related datasets (AEs, CM, labs)
- ✓ **Reduce review burden:** Shorten time to database lock through intelligent prioritization

 **Human role:** Final query decisions and database lock



Statistical Programming

- ✓ **Draft specifications:** Auto-generate SDTM and ADaM specs from protocol & eCRF
- ✓ **Generate code:** Draft transformation code and TLF templates for standard analysis
- ✓ **Validation checks:** Suggest derivation logic and validation checks for complex endpoints

 **Human role:** Validate complex derivations and scientific interpretation



Regulatory Submissions

- ✓ **Auto-populate documents:** Pre-fill submission sections from source data and prior study text
- ✓ **Consistency checks:** Verify alignment across protocols, SAPs, CSRs, and dossiers
- ✓ **Draft responses:** Assist with early drafting of regulatory agency responses

 **Human role:** Robust review and QC before finalization

Current Benefits & Limitations

A realistic view of AI's impact in clinical development today



Benefits

Reduced Repetitive Workload

Significant reduction in manual tasks like drafting, coding, and data review, freeing up valuable time

Improved First-Pass Quality

AI-generated drafts and checks are more consistent and complete, reducing revision cycles

Faster Access to Insights

Accelerated preliminary analyses enable faster decision-making during critical trial milestones

Reallocation to High-Value Work

Teams shift focus from rote execution to complex scientific interpretation and strategic analysis



Limitations

Hallucination Risks

LLMs can generate inaccurate or fabricated content, requiring rigorous human verification

Inconsistency Issues

AI outputs may lack consistency across different runs or contexts, posing standardization challenges

Integration Challenges

Embedding AI into validated workflows involves complex technical and compliance hurdles

Evolving Regulatory Expectations

Regulatory guidance is still developing, creating uncertainty and requiring proactive governance



Today's Reality

AI is a useful assistant and co-pilot under human control and review — it augments productivity but does not replace human oversight. **Validation and governance frameworks are essential.**

AI Maturity Model: From Assistive to Agentic

Three stages of AI capability representing increasing autonomy and proactivity

1

PAST 2+ Yrs

Assistive AI

Reactive AI

Responds to user prompts; drafts content or suggests queries and checks

Human: All execution and decision-making

2

NEAR-TERM

Co-Pilot AI

Orchestrated AI

Triggers predefined workflows and orchestrates multiple steps once initiated

Human: Initiates and owns decisions

3

Next 3-5 YEARS

Agentic AI

Proactive AI Agents

Monitors continuously and initiates tasks when conditions are met; escalates to humans

Human: Supervises and intervenes when needed



The Transition Journey

This paper focuses on the transition from today's assistive/co-pilot systems to emerging **agentic systems over the next 3-5 years**. Each stage builds upon the previous, with increasing autonomy bounded by human oversight and regulatory compliance.

Latest Agentic AI Developments

Agentic AI has moved from experimental curiosity to business necessity

 **\$199B**
by 2034

Market growing from **\$5.25B (2024)** at **43.84% CAGR**

Source: Arcade.dev, 2025

 **79%**
Adoption Rate

of organizations report AI agent implementation

Critical mass achieved

 **171%**
Average ROI

U.S. enterprises achieving **192% ROI**

Exceeded expectations

 **Key Enterprise Metrics**

 **Enterprise Apps with AI Agents**
By end of 2026 (Gartner)

40%

 **Multi-Agent Architectures**
Of implementations

66.4%

 **AI Budget to Agentic**
Companies allocating >50%

43%

 **Top Deployment Challenges**

 **Governance & Security**
75% of tech leaders cite as primary challenge

 **Project Cancellation Risk**
>40% of projects canceled by 2027 (Gartner)

 **Agent Reliability**
Unrealistic expectations about capabilities

 **Biotech Implication:** The rapid enterprise adoption validates agentic AI's potential—regulated industries must prepare governance frameworks now to capture value safely.

GitHub Copilot's Evolution to Autonomous Agent

From code completion to autonomous coding partner—lessons for clinical programming

The Journey

1

2021: Code Completion

Simple autocomplete suggestions in IDE

2

2023: Chat & Explain

Interactive assistance and code explanation

3

2026: Autonomous Agent

Full task execution with human oversight

20M+

Developers

90%

Fortune 100

MCP: Model Context Protocol

Standardized protocol enabling agents to securely access external tools and data sources. **97M+ monthly SDK downloads**. Allows Copilot to integrate with third-party systems while maintaining security.

New Agent Capabilities (2026)



Autonomous Refactoring

Multi-file code restructuring



Bug Detection & Fixing

Identifies and resolves issues proactively



Test Generation

Creates comprehensive test suites



Pull Request Creation

Full GitHub workflow integration



Biotech Application

Clinical Programming Parallel: Progress from AI-assisted SDTM/ADaM specification drafting → autonomous dataset generation → continuous pipeline execution with validation checks and human approval gates.

“ EY DevEx Lead: "The GitHub Copilot coding agent opens doors for each developer to have their own team, all working in parallel to amplify their work." — James Zabinski, 2025

DevOps Automation & CI/CD Practices

Continuous integration and deployment—lessons for clinical data pipelines



38%

Deployment pipeline time reduction through AI optimization



20-40%

Automation in routine DevOps tasks achieved



Proactive

Issue detection before production deployment



Core CI/CD Practices

1

Continuous Integration
Automated code integration and testing

2

Continuous Deployment
Automated release to production

3

Infrastructure as Code
Version-controlled environment setup

4

Automated Testing
Self-healing test suites



AI Agent Achievements

✓

Pipeline Optimization
AI rewrites slowest stages, refactors caching rules, suggests better branch policies

✓

Pre-Production Detection
Detects misconfigured permissions before push—"basically our third SRE"

✓

Self-Healing Systems
Auto-remediation with guardrails: restart crashed workloads, scale hot paths



Biotech Application: Clinical Data Pipelines

Adopt CI/CD Principles:

- Auto-rerun analyses when new data arrives
- Continuous validation of SDTM/ADaM outputs
- Automated QC checks with human approval gates

Pipeline Automation:

- Event-driven triggers for TLF generation
- Version-controlled specifications and code
- Rollback capabilities for validated outputs

Site Reliability Engineering (SRE) Practices

Observability and proactive monitoring—lessons for clinical trial oversight



Logs

Detailed event records for analysis and troubleshooting



Metrics

Numerical performance indicators and thresholds



Traces

End-to-end request flow visualization

Key SRE Practices



Observability

OpenTelemetry for unified telemetry collection



Proactive Monitoring

Automated alerting on actionable thresholds



Error Budgets & SLOs

Service Level Objectives with measurable targets



Auto-Remediation

Guardrailed self-healing with rate limits

SRE Tool Stack

Prometheus

Metrics

Grafana

Dashboards

OpenTelemetry

Instrumentation

PagerDuty

Incident Mgmt

Jaeger/Tempo

Distributed Tracing



Biotech Application: Clinical Trial Monitoring

Implement SRE Principles:

- Continuous trial monitoring with automated safety signal detection
- Predictive enrollment forecasting using historical patterns
- Management-by-exception workflows for data managers

Observability for Trials:

- Real-time data quality metrics and anomaly detection
- Site performance dashboards with drill-down capabilities
- Automated escalation when thresholds breached

Agentic AI in 3-5 Years: Near-Future Capabilities

Proactive, bounded autonomy systems that amplify human reach



Continuous Monitoring

Real-time surveillance of trial data, safety signals, and operational metrics



Auto-Execute QC

Automatic execution of quality checks when data thresholds are crossed



Pipeline Automation

Auto generation of SDTM/ADaM/TLF pipelines when new data arrives



Predictive Analytics

Forecast enrollment, operational risks, and trial outcomes proactively



Digital Twins

Run simulations to model trial scenarios and optimize designs



Auto-Assembly

Assemble draft reports and submission documents as data accumulate

↔ The Transformation Impact

Batch



Continuous

Trial execution shifts from episodic to always-on

Weeks



Hours

Time to insight dramatically compressed

Execute



Supervise

Roles evolve to oversight and interpretation

Reactive



Proactive


Quality management becomes preventive



Key Principle: Agentic AI does not replace scientific accountability—it amplifies human reach and enables management by exception.

AI Maturity Model: From Assistive to Agentic

Three stages of AI capability representing increasing autonomy and proactivity



1 Assistive AI

Reactive AI

Responds to user prompts; drafts content or suggests queries and checks

Human: All execution and decision-making

2 Co-Pilot AI

Orchestrated AI

Triggers predefined workflows and orchestrates multiple steps once initiated

Human: Initiates and owns decisions

3 Agentic AI


Proactive AI Agents

Monitors continuously and initiates tasks when conditions are met; escalates to humans

Human: Supervises and intervenes when needed

24-Month Implementation Roadmap

Start small, scale gradually, with governance and measurement throughout



FDA-Aligned Agentic AI Framework

Six pillars aligned with FDA's 10 guiding principles for safe, compliant deployment



FDA: 1, 4, 10

Human-Centric Design

Support—not replace—clinical decision-makers with clear context of use, boundaries, and limitations



FDA: 2, 3, 4, 5

Risk-Based Governance

Risk-based approach determines validation depth, sign-off requirements, and change controls



FDA: 6, 7

Data & Engineering

Standardized, well-governed data with documented lineage and good engineering practices



FDA: 1, 8

Bounded Autonomy

Agents operate within defined boundaries: observe → act → escalate with human-in-the-loop



FDA: 2, 8, 9

Validation & Monitoring

Formal risk-proportionate validation before deployment with ongoing monitoring for drift



FDA: 6, 7, 10

Transparency & Auditability

Explainable, traceable behavior with logs and versioning for inspection reconstruction



Outcome

When appropriately governed and validated, agentic AI enhances evidence generation, reduces operational latency, and expands human oversight capacity in clinical development.

Governance & Validation Requirements

Essential elements for safe, compliant agentic AI deployment in GxP environments



Pre-Deployment Validation

Risk-proportionate validation must be completed before any agentic system goes live, including performance testing, edge case evaluation, and human-AI interaction assessment



Ongoing Monitoring

Continuous surveillance for model drift, safety signals, and compliance risks with automated alerts and periodic performance reviews



Explainability & Traceability

Agent decisions must be interpretable with complete audit trails that allow reconstruction of actions during inspections



Cross-Functional Involvement

Clinical, statistics, data management, QA, regulatory, and AI engineering must collaborate throughout the lifecycle



Clear Communication

Plain-language documentation of agent capabilities, limitations, and appropriate use cases for all stakeholders



Change Control & CAPA

Formal change control processes and corrective/preventive action (CAPA) workflows to maintain reliability across the agent lifecycle



Central Principle

Human expertise remains central at every stage—AI augments capabilities but never replaces scientific accountability and regulatory responsibility.

The Path Forward: Embracing Pragmatic Evolution

From episodic batch processes to continuous data-driven cycles

↔ The Operating Model Shift



🎯 Key Outcomes

- Shorter Time-to-Insight**
Rapid access to trial intelligence accelerates decision-making
- Proactive Quality Management**
Issues detected and addressed before they become problems
- Increased Scientific Agility**
Teams focus on interpretation rather than rote execution

💡 The Approach

Start with assistive use cases, mature into orchestrated co-pilot automations, and introduce agentic behaviors **only when data, governance, and validation are ready**.



The Ultimate Goal

By investing in safe foundations and embracing pragmatic evolution, BioPharma organizations can expand oversight capacity, accelerate evidence generation, and ultimately **improve patient access to new therapies**.