Paper AS02

Revolutionizing Clinical Quality: A Frontier Approach with Quality by Design (QbD), Infrastructure Modernization and Data-Driven Analytics

Priti Gupta, GSK, Bangalore, India
Dr. Deepa Balaji, GSK, Bangalore, India

ABSTRACT

Clinical research has reached a critical juncture. Trials that once relied on paper case report forms and relatively simple protocols now leverage adaptive designs, decentralized components and multiple digital endpoints. Data streams from central laboratories, imaging vendors, wearables and home health devices are ingested at near real-time cadence into sponsor environments.

Modern clinical trials commonly generate more than 3.5 million data points in Phase III studies. Despite this data abundance, many organizations still rely on legacy quality-assurance practices involving manual source data verification (consuming 30-40% of budgets), uniform site monitoring approaches and largely reactive issue detection. This whitepaper presents transformation framework built on three core pillars—**Quality by Design (QbD), Infrastructure Modernization and Three-Layer Data-Driven Analytics**—to enable proactive, risk-proportionate and automated oversight.

INTRODUCTION

Clinical research is undergoing a profound transformation driven by:

1. **Rising Data Volume and Velocity**: eSource, ePRO/eCOA, wearable sensors, advanced imaging, real-world evidence and decentralized trials generate unprecedented data at scale, demanding new quality assurance approaches.

2. **Increasing Protocol Complexity**: Adaptive trial designs and composite endpoints raise the bar for statistical rigor and operational consistency, introducing significant challenges in maintaining reliability.

3. **Evolving Regulatory Expectations**: Global regulators now require proactive, risk-based quality management strategies with robust computer system assurance and traceable data governance.

4. **Resource Constraints**: Traditional monitoring methods and manual reconciliation are unsustainable, cannot scale effectively and are increasingly cost-inefficient.

Key benefits of implementation documented across diverse trial contexts include: - Large Phase III trials: 25-40% cost reduction, 50-60% monitoring reduction - Phase I/small-scale studies: 20-35% cost reduction, 40-50% monitoring reduction - Future-state digital twin implementations: 30-40% cost reduction, 60-70% monitoring reduction, 8–12-week earlier risk detection

The framework provides a structured, phased adoption pathway enabling organizations to incrementally mature their quality infrastructure while capturing value at each implementation stage.

1. The Three Foundational Pillars

Quality in clinical trials must be engineered rather than inspected in, automated rather than manual and predictive rather than reactive. The three-pillar framework integrates:

Pillar 1: Quality by Design (QbD) – Engineering quality into trial design from day 1 through systematic identification of Critical-to-Quality (CtQ) factors, formal risk assessment and proportionate monitoring strategies.

Pillar 2: Infrastructure Modernization – Deploying unified, cloud-driven, analytics-ready data ecosystems that replace fragmented workflows with real-time data synchronization, sub-24-hour latency and automated quality controls.

Pillar 3: The Three-Layer Data Driven Quality Oversight Model – Operating across AI/ML based Intelligent query detection, supervised Key Risk Indicators (KRIs) with Quality Tolerance Limits (QTLs) and unsupervised data anomaly detection.

Together, these three pillars create a closed-loop, risk-proportionate quality system that shortens cycle times, reduces regulatory findings, enhances data integrity, provides "fit for purpose data" and ultimately improves patient safety—while maintaining full alignment with ICH E6(R3) and FDA AI governance frameworks.

1.1 Quality by Design and Critical-to-Quality Factors

Quality by Design (QbD) represents a paradigm shift from retrospective inspection to proactive engineering of quality. Rather than assuming all collected data are equally important, QbD questions as to which aspects of this trial are truly critical to patient safety and to drawing valid, reliable conclusions.

QbD ensures that endpoints reflect meaningful patient outcomes, improving regulatory acceptance and patient engagement. Protocols are simplified by removing non-essential procedures, reducing site burden and data noise. Risks are identified upfront with specific mitigations tied to measurable Key Risk Indicators. Monitoring is proportionate, allocating resources based on risk tier rather than uniform 100% source data verification.

The step-wise QbD Implementation Process

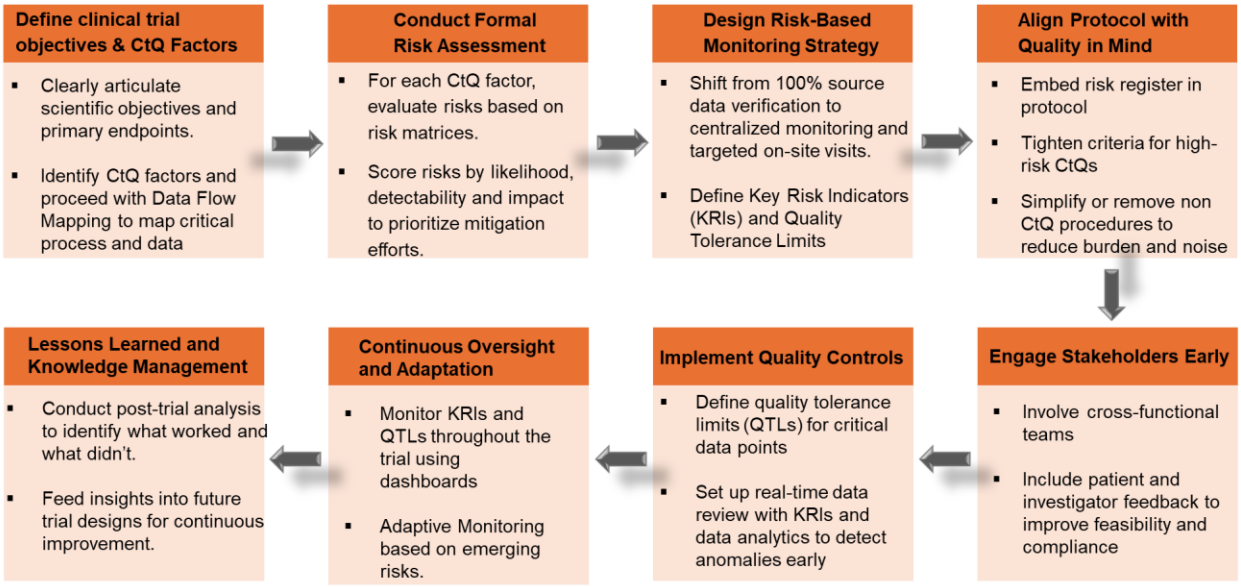

Before vs. After: QbD Impact

Clinical Trial Aspect	Before QbD (Traditional)	After QbD (with CtQ Focus)
Protocol Complexity	Many "just in case" procedures creating burden	Simplified protocols, non-CtQ procedures removed
Monitoring Focus	Broad monitoring of all variables equally	Targeted monitoring linked to CtQ factors
Risk Handling	Issues discovered reactively via deviations/audits	Risks identified upfront with specific mitigations
Protocol Amendments	Higher rates (3.2 per study average)	Reduced (1.8 per study with QbD)
Site Experience	Quality perceived as compliance burden	Quality seen as shared responsibility reducing overall burden
Database Lock Timeline	52 weeks typical	46 weeks (6-week acceleration)
Regulatory Findings	2-3 483 forms on an average	Zero findings with full transparency

1.2 Infrastructure Modernization

Legacy clinical architecture (2010-2018) relied on fragmented, manual workflows. These legacy systems were characterized by dual-entry burdens (paper + EDC), high dependency on on-site monitoring (50-100+ visits per trial) and latency-heavy data flows that took 2-4 weeks to move from site to programming. These bottlenecks delayed signal detection and inflated operational costs. In contrast, the Modern Architecture (2024+) utilizes a unified, cloud-driven ecosystem with near-continuous synchronization. By integrating eSource and wearables directly into the cloud, organizations have shifted to a “default-remote” monitoring model. This architecture enables <24-hour data latency, allowing for near real-time dashboards and proactive quality signals that identify risks before they escalate.

The Data Shared Service (DSS) Core acts as the centralized technical engine, functioning through two primary mechanisms:

- **Centralized Data Acquisition**: DSS synchronizes diverse streams—including EDC (via FTPS/sFTP), CTMS, Lab, ePRO and RWE—into a single source of truth.

- **Standardized Pipelines**: Data is processed via SAS or Python scripts to harmonize identifiers and align with standards. These API-driven pipelines create application-specific, analytics-ready snapshots that are fully auditable and reusable.

Integrated Infrastructure and Security: Modernized infrastructure transforms legacy systems into near real-time intelligence hubs by aggregating data from key repositories like RAMOS (IVRS data), CDB, EDC, Vendor and CTMS. This ecosystem is supported by three technological pillars:

- **Cloud Scalability**: GxP-compliant AWS/Azure environments provide the high availability required for global trials.

- **Hybrid Ingestion Layer**: The system supports both legacy and modern data transfers, utilizing secure sFTP drop boxes, FTPS and a new Data Exchange API for direct, near real-time ingestion.

- **Interoperability and Transparency**: By integrating comprehensive audit trail data from all sources, the system ensures full data lineage.

Modern infrastructure must meet the highest regulatory and cybersecurity standards: 21 CFR Part 11 for electronic records and signatures, GDPR/CCPA for global data privacy and residency requirements, enterprise cybersecurity including encryption and intrusion detection, complete audit trails of data access and transformations and role-based access control (RBAC) ensuring principle-of-least-privilege.

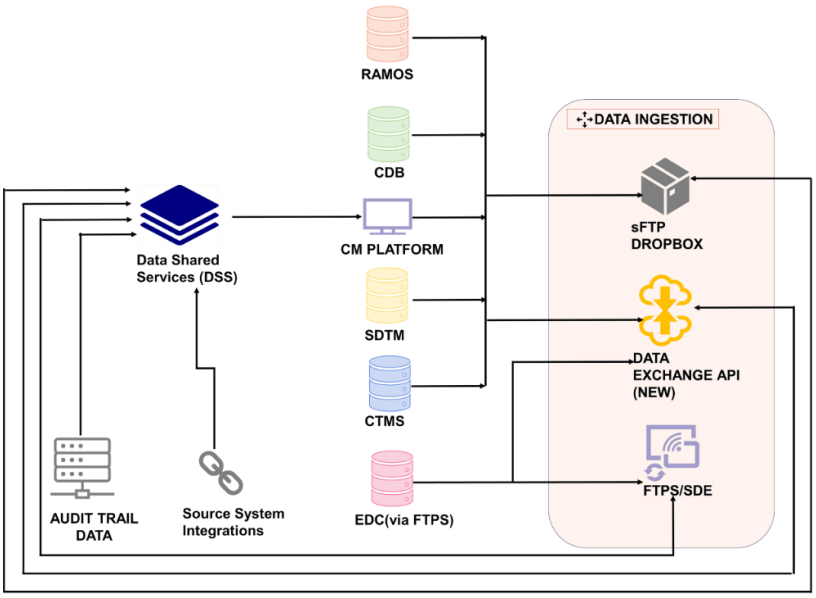

Figure: Enablers of Infrastructure Modernization

1.3 Analytics as an Enabler of Proportionate Oversight

Analytics provides an operational engine that transforms QbD insights and modern infrastructure into actionable oversight. However, analytics must be:

- **Proportionate**: Focused on CtQ-relevant risks rather than generating noise across all variables

- **Transparent**: Interpretable by clinicians, statisticians and quality professionals to support defensible decisions

- **Integrative**: Embedded in workflows so that alerts, trend shifts and anomalies are reviewed and acted upon

2. The Three-Layer Quality Oversight Model

The three-layer oversight model is a progressive framework designed to modernize clinical quality management by shifting from manual, rule-based processes to an AI-driven, proactive strategy. This model improves precision, reduces false positives by up to 95% and allows experts to focus on complex risks rather than manual data cleaning.

2.1 Layer 1: Intelligent Query Detection

Layer 1 focuses on discrete, record-level data inconsistencies and protocol deviations using rule-based or rule-plus-pattern logic applied at or near the time of data entry.

Typical categories include - Range violations: Physiologically implausible laboratory values, Logical inconsistencies: SAE flag set to “Yes” but severity recorded as “Mild” with no hospitalization, Temporal sequence errors: Dosing before informed consent, procedures outside visit windows, Direct protocol deviations: Expressed algorithmically.

Detected issues trigger automated queries constructed from preapproved templates that reference relevant data elements, provide clear clinical context and specify required site responses.

Benefits of Layer 1 implementation: - Reduction in manual review effort for data managers by 50-60%. Decreased query cycle times from weeks to days. Improved consistency and traceability in issue identification and resolution

2.2 Layer 2: Key Risk Indicators and Quality Tolerance Limits

Layer 2 focuses on aggregated trial performance metrics directly linked to CtQ factors, operating at site, country and study levels to surface emerging trends before they become critical. By using the “cleaned” data from Layer 1, these metrics become more reliable indicators of systemic issues.

Key Risk Indicators (KRIs) monitor performance through time-series metrics: - Timeliness of adverse event reporting - Visit adherence and missed visit rates - Data entry lag (e.g., median days from visit to data entry) - Protocol deviation rates and types - Enrollment rates relative to expectations

Quality Tolerance Limits (QTLs) are higher-level thresholds that, if breached, indicate a serious systemic risk requiring documented investigation and potential regulatory disclosure. Examples include: - Cumulative rate of major eligibility violations - Frequency of important dosing errors - Systemic underreporting of SAEs

Interactive dashboards, statistical process control charts and drilldown capabilities allow cross-functional teams to interpret KRI and QTL behavior. When risk signals cross predefined thresholds, study teams can retarget onsite or remote monitoring visits, update training, initiate root cause analyses or adjust KRIs as new information emerges.

Effective KRI-driven oversight has been associated with 50-60% reductions in routine on-site monitoring visits while maintaining or improving quality.

2.3 Layer 3: Data Quality Assessment and Expert Manual Review

Layer 3 addresses cross-domain inconsistencies and professional participant detection not easily resolved through static rules or simple aggregations. Key components include:

1. Systematic assessment of completeness and consistency across domains (verifying that each serious adverse event has corresponding narratives, lab follow-up and concomitant medication records where clinically appropriate)

2. Duplicate or professional participant detection, particularly important in Phase I studies or geographically dispersed multisite programs.
3. Structured expert manual review, where experienced clinicians, statisticians, and data scientists examine patient-level or site-level profiles to interpret emerging patterns, suppress false positives, and contextualize anomalies

Layer 3 is **"human-in-the-loop"** approach ensures that expert time is spent on high-value adjudication rather than routine screening.

3. The Three-Layer Quality Oversight Model in Practice with QbD Alignment

3.1 Case Study 1: Large Phase III Oncology Trial

A Phase III oncology trial enrolled approximately 300 participants across 40 global sites. CtQ factors included reliable imaging-based assessment of tumor response, accurate and timely SAE reporting, correct dosing and drug accountability

Framework Implementation:

- Layer 1 rules checked for dosing outside protocol-defined windows, missing SAE seriousness attributes, inconsistent imaging dates and mismatched progression dates versus RECIST assessments

- Layer 2 tracked KRIs such as SAE reporting timeliness, imaging schedule adherence and visit completion rates by site and country. QTLs set for missing or "not evaluable" RECIST assessments over time and cumulative rates of major dosing deviations.

- Layer 3 involved quarterly multidisciplinary review cycles examining integrated patient profiles and duplicate detection algorithms screening for anomalously similar demographics and visit patterns

Key Outcome:

1. **During Layer 3 review**, duplicate patient detection flagged two participants enrolled at different sites whose demographics, baseline laboratory patterns and visit timelines were highly similar. Manual investigation confirmed the same individual had been inadvertently enrolled twice. One record was excluded, corrective actions and implemented.

2. **Early signal detection and improvement in PFS data**: Layer 1 detected scans performed outside the protocol visit windows, RECIST inconsistency flags: progression date captured earlier than the central read's PD confirmation; Layer 2 with QTL detected Not Evaluable rate for RECIST at rolling 4-week, the NE proportion exceeded threshold; corrective imaging and charter reinforcement reduced NE to baseline.

3. **SAE reporting timeliness** – EDC checks detected initial SAE entries with missing hospitalization dates, causality and action taken. A small cluster of sites exhibited longer median time from awareness to initial report which was escalated to Operational team for targeted retraining.

Estimated cost avoidance was approximately USD 150,000-200,000 in avoided monitoring, reanalysis and remediation work.

3.2 Case Study 2: Phase I Rare Disease Study

Small-scale studies present distinct challenges: limited statistical power for comparative metrics, high duplicate enrollment risk, resource constraints and manual data entry dominance. A Phase I dose-escalation study enrolled 24 participants with a rare metabolic disorder across three specialist centers

Adapted Framework:

- Layer 1 employed enhanced query logic for manual data entry, checking for missing informed consent dates, dosing outside cohort windows and temporal inconsistencies

-Layer 2 shifted to patient-level KRIs (dosing adherence, visit punctuality, AE reporting lag, safety assessment completion) with threshold-based QTLs rather than comparative site benchmarks

- Layer 3 centered on structured patient-level review cycles conducted weekly during dose escalation and biweekly during dose expansion

Key Outcomes:

1. **Eligibility Flag Raised**: Week 3 flagged a participant who had previously enrolled in a competitor's study, preventing enrollment before dosing commenced

2. **Early safety signal detection**: Week 5 identified a cluster of elevated transaminases in 3 of 4 participants in Cohort 2, prompting intensified monitoring and preventing a missed safety signal

3. **Data quality improvement**: Week 8 detected missing dosing documentation and resolved it through targeted retraining

Cost avoidance relative to traditional weekly on-site visits: approximately USD 40,000-60,000 over the study duration.

4. Advanced Analytics and AI/ML Integration

4.1 Digital Twins and Contextual Anomaly Detection

Digital Twins

A digital twin is a dynamic, computational replica of a real-world system (a trial site or patient cohort) that continuously ingests observational data, incorporates domain knowledge and allows simulation of scenarios and prediction of outcomes.

In clinical quality, digital twins serve to:

- Synthetic Control Arms: Virtual models act as control groups, reducing or replacing real placebo arms. Minimizes patient burden and trial size while maintaining statistical power.

- Outcome Prediction & Treatment Effect Estimation: Patient-specific digital twins simulate standard-of-care responses. Enables precise estimation of treatment effects and improves trial sensitivity.

Optimized Trial Design: Simulates trial designs and protocols (e.g., dosage regimens, inclusion criteria, visit schedules). Helps refine design, reduce sample size, and minimize patient dropout.

- Personalized Dosing & Biomarker Assessment: Predicts optimal dosing for individuals or subgroups. Identifies early biomarkers for efficacy and safety—critical for AI-designed drug studies.

- Real-time Adverse Event Forecasting: Continuously updated digital twins monitor participants. Foresees potential safety issues, improving patient protection and trial monitoring.

- Contextualize anomalies: When an anomaly is detected (e.g., zero AEs at a high-enrollment site), the twin can simulate plausible operational explanations

- Logistics & Operational Optimization: Simulates supply chain logistics (enrollment pacing, site capacity, CRA workload, cold chain integrity). Enhances resource allocation and reduces trial delays.

Contextual Anomaly Detection (CAD)

Contextual Anomaly Detection (CAD) identifies data points that appear "normal" alone but are unusual given trial context (site, patient profile, visit timing, protocol expectations). They can signal protocol deviations, atypical site behavior or inconsistent patient data—supporting earlier risk detection for data integrity, patient safety and compliance.

Symptom clusters refer to groups of related symptoms that tend to occur together in patients during a clinical trial. Instead of looking at individual symptoms in isolation, clustering helps identify patterns—such as a combination of headache, nausea and fatigue appearing frequently in the same group of participants may indicate emerging adverse effects (e.g., headache+nausea+rash → hypersensitivity; fatigue+shortness of breath → potential cardiotoxicity).

Workflow: collect symptom data → cluster patients (e.g., k-means/hierarchical) → assess frequency/severity and association with AEs → flag clusters that exceed thresholds for medical review.

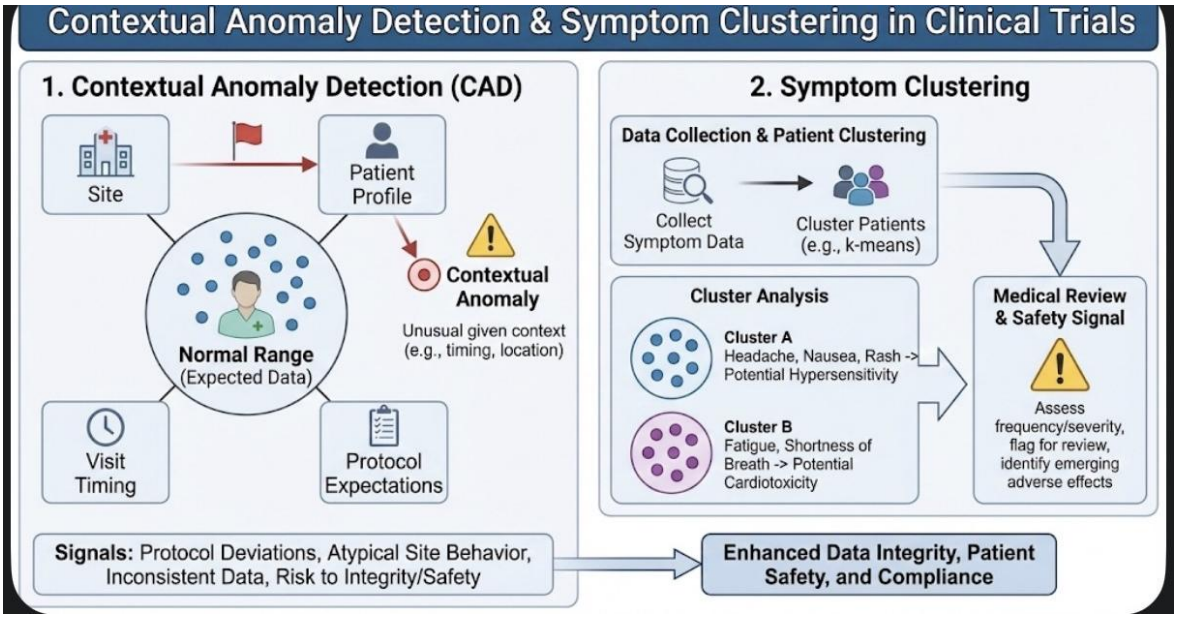

Source: Gemini Pro

Synthetic Data and NLP: Synthetic data enables training and validation of anomaly detection models without exposing real participant data, addressing data governance and regulatory concerns. Natural language processing (NLP) can extract structured insights from unstructured AE narratives and source documents, identifying semantic anomalies or missing clinical context.

4.2 Future-State Implementation: Dose Optimization + Digital Twin Augmentation to Reduce Sample Size and Accelerate Timelines

To demonstrate that future-state clinical trial execution can deliver measurable efficiency gains, the below case study presents a simplified and practical implementation of dose optimization and digital twin augmentation in a mid-sized interventional study. The primary objective is to reduce randomized sample size and shorten the overall trial timeline while maintaining or improving statistical power and ensuring robust data quality.

Traditional design.

A sponsor planned a Phase II study for a novel oral therapy in Type 2 Diabetes with a 12-week endpoint assessing change in HbA1c. The original design was a conventional three-arm randomized controlled trial (Placebo vs Dose A vs Dose B) with the following assumptions:

• Sample size: 360 participants (120 per arm)

• Sites: ~45 sites

• Primary endpoint: HbA1c reduction at Week 12

• Expected duration: ~12 months (start-up, enrollment, follow-up, database lock)

While scientifically sound, the traditional approach carried predictable inefficiencies such as a high probability of exposing participants to suboptimal dosing, increased discontinuations due to tolerability differences between doses and longer timelines driven by higher enrollment and operational burden.

Future-state implementation approach.

The sponsor redesigned the study using a two-stage approach incorporating QbD principles, model-informed dose selection, and controlled digital twin augmentation:

1. Stage 1: Dose optimization lead-in (small cohort). A lead-in cohort of 40 participants was enrolled to evaluate early safety, tolerability and preliminary HbA1c response across Dose A and Dose B. Based on observed discontinuations and early response trends, Dose A was selected as the optimal dose to carry forward.

2. Stage 2: Confirmatory trial with reduced placebo allocation. Instead of continuing with a full three-arm design, the confirmatory stage was executed as a two-arm trial (Placebo vs Optimal Dose) with reduced placebo enrollment: Placebo: 80 participants; Optimal dose (Dose A): 160 participants; Total randomized: 240 participants. In addition, the sponsor implemented a digital twin augmentation strategy that contributed placebo-equivalent information corresponding to approximately 40 additional placebo participants, subject to prespecified similarity and concordance diagnostics and supported by sensitivity analyses to ensure robustness.

3. A three-stack analytical model and CAD-enabled monitoring. A three-stack analytical model was used to improve data reliability and shorten database lock cycles. CAD was applied to identify anomalies and trends in near real time (e.g., discontinuation clusters, endpoint distribution shifts, site-level outliers), enabling earlier corrective actions and reducing downstream query volume.

Quantified outcomes and benefits.
The redesigned future-state implementation achieved measurable improvements compared with the traditional plan:

- Randomized sample size reduction: from 360 to 240 participants, representing a 33% reduction (120 fewer participants).
- Increased statistical power: from 80% to ~88–90%, driven by optimized dose selection (lower discontinuations) and increased effective placebo information through controlled augmentation.
- Faster timeline: from approximately 12 months to 9 months, representing a 25% reduction in overall duration (≈3 months faster).
- Lower dropout rate: from 18% to 10%, representing a 44% reduction, primarily due to improved tolerability and proactive monitoring.
- Cost avoidance: assuming an average cost of $25,000 per participant, enrollment reduction generated ~$3.0M in direct savings. Additional operational efficiencies (reduced monitoring burden, fewer queries, and faster database lock) contributed an estimated $0.5M–$1.0M, resulting in ~$3.5M–$4.0M total savings.

Key takeaway
This simplified case study illustrates that a practical combination of dose optimization, controlled digital twin augmentation and A three stack analytical model can reduce enrollment requirements, increase statistical power and accelerate time-to-decision. Importantly, the approach supports QbD principles by proactively managing risk (dose tolerability, data variability and site performance) while preserving scientific rigor and improving overall trial efficiency.

5. EMA–FDA Guiding Principles for Good AI Practice and Implementation Roadmap

Regulators increasingly recognize that Artificial Intelligence (AI) and Machine Learning (ML) can improve the efficiency, sensitivity, and scalability of evidence generation—provided these technologies are deployed with robust governance, transparency, and lifecycle controls.

The EMA and the U.S. FDA jointly describe guiding principles for “Good AI Practice” in drug development, intended to ensure that AI-enabled systems remain reliable, explainable, secure, and fit-for-purpose across the product lifecycle. In clinical quality oversight, these principles translate into practical requirements for AI-enabled RBQM, centralized monitoring, and data-driven risk detection.

This section operationalizes the principles into clinical quality deliverables, enabling inspection-ready adoption while maintaining human accountability and CtQ-driven oversight.

5.1 EMA–FDA Guiding Principles and Practical Translation for Clinical Quality Oversight

#	Guiding Principle	What it requires (Regulatory intent)
1	Human-centric by design	AI must support ethical, human-centered decision making
2	Risk-based approach	Oversight and validation proportional to intended use and risk
3	Adherence to standards	Alignment with GxP expectations, cybersecurity, privacy, and scientific rigor
4	Clear context of use (CoU)	Explicit definition of what the AI does and does not do
5	Multidisciplinary expertise	Continuous involvement of domain experts
6	Data governance & documentation	Provenance, integrity, bias awareness, privacy protection
7	Model design & development practices	Reproducibility, explainability, and controlled development
8	Risk-based performance assessment	Validate the full system including human-AI workflow
9	Lifecycle management	Ongoing monitoring, change control, periodic reassessment
10	Clear, essential information	Users must understand limitations, updates, and outputs

Applications of these principles in Clinical Trials

These principles directly impact trial operations across multiple domains but here we will discuss Risk-Based Quality Management (RBQM) & Anomaly Detection.

AI-Driven Monitoring for Data Integrity & Patient Safety

Principle 2 (Risk-Based Approach): Prioritize AI oversight for high-risk data (safety labs, vital signs, concomitant medications) and clinical decision points. Use machine learning for contextual anomaly detection—flagging deviations from expected patterns based on patient demographics, disease status, and site practices.

Principle 5 (Multidisciplinary Expertise): Collaborate with clinical monitors, data managers, biostatisticians, and IT security to design AI models that integrate with existing RBQM frameworks (e.g., Central Tendency Monitoring, Aggregate Analysis).

Principle 9 (Lifecycle Management): Continuously monitor AI model performance —retrain quarterly or when new patient cohorts enroll. Adjust alert thresholds if site patterns evolve (e.g., new monitoring equipment, staff turnover).

Principle 6 (Data Governance & Documentation): Maintain detailed logs of AI flagged issues, clinical monitor investigations, and resolutions. Use these for regulatory inspection readiness and continuous improvement.

Principle 1 (Human-Centric by Design): AI flags issues and recommends actions (e.g., "review this patient's glucose logs"), but clinical monitors retain decision authority. Ensure workflows keep humans in the loop.

Impact: Faster detection of data quality issues, improved patient safety through early identification of adverse trends, reduced manual monitoring burden (20-40% efficiency gains), enhanced regulatory compliance.

5.2 Implementation Roadmap: Phased Adoption in Clinical Quality Oversight

AI-enabled clinical quality oversight should be implemented in phases, allowing organizations to capture value early (e.g., reduced query burden, faster issue detection) while maturing governance, validation, and lifecycle monitoring.

Phase	Typical timeline	Primary objective	Key activities (What to implement)	Key deliverables (What to show)
Phase 0 – Define & govern	0–3 months	Establish safe and compliant foundation	Define CoU (Context of use), decision rights, governance, model risk tiering; align AI outputs to CtQs and RBQM strategy	CoU; RACI; AI risk assessment; CtQ/KRI/QTL linkage; privacy approvals
Phase 1 – Build Foundations	3–9 months	Enable reliable, auditable AI inputs	Standardize data pipelines; enforce lineage + auditability; implement baseline transparency (model cards); strengthen CSA	Dataset specs; lineage maps; Data Quality checks; CSA (Computer Software Assurance) artifacts (fit for purpose software); documentation templates
Phase 2 – Pilot & validate	9–18 months	Prove performance in controlled trials	Run pilots on selected studies; validate end-to-end workflows; calibrate thresholds; define escalation and oversight SOPs	Validation report; workflow evidence; usability testing; SOPs; monitoring plan
Phase 3 – Scale & operationalize	18–36 months	Embed AI into RBQM operations	Expand across portfolio; automate drift monitoring; integrate triage into QMS; strengthen training and change control	Dashboards; drift metrics; CAPA process; training records; periodic re-validation plan
Phase 4 – Optimize & Innovate	36+ months	Mature to predictive intelligence	Implement contextual anomaly detection and digital twin augmentation where justified; engage regulators early	Advanced governance pack; inspection/submission readiness package; continuous improvement evidence

6. Conclusion: Advancing the Next Generation of Clinical Quality

Clinical quality oversight is entering a decisive new era. The increasing scale, complexity, and digital intensity of clinical trials have exposed the limitations of traditional, inspection-driven quality models that rely on uniform monitoring, heavy manual effort, and retrospective issue detection. To remain effective and credible, quality must transition from a reactive compliance function to a **proactively designed, risk-proportionate system** embedded throughout the trial lifecycle.

This paper presented a practical and future-ready framework that integrates **Quality by Design (QbD)**, **modernized data infrastructure**, and a **three-layer, data-driven quality oversight model**. By anchoring oversight to **Critical-to-Quality (CtQ)** factors, the approach ensures that quality activities are focused on what truly impacts patient safety and the reliability of clinical conclusions—rather than exhaustive verification of all data points.

Modern, cloud-enabled infrastructure acts as the foundation for this shift, enabling near real-time data visibility, full traceability, and inspection-ready auditability. Building on this foundation, the three-layer oversight model **systematically embeds quality at scale**: intelligent query detection strengthens baseline data reliability; KRI- and QTL-driven monitoring enables early identification of emerging systemic risks; and structured expert review ensures that complex anomalies are interpreted within appropriate clinical and operational context.

As organizations mature along this pathway, advanced analytics—such as contextual anomaly detection, synthetic data strategies, and digital twin augmentation—offer additional opportunities to accelerate signal detection, optimize trial design, and reduce unnecessary monitoring and participant exposure. When implemented through a phased, regulator-aligned roadmap with clear governance and human oversight, these capabilities can enhance both efficiency and regulatory confidence.

Ultimately, the future of clinical quality will be defined not by technology alone, but by the **disciplined integration of**

risk-based quality principles, fit-for-purpose data ecosystems, and expert human judgment augmented by intelligent systems. Organizations that adopt this model will be better positioned to deliver faster, safer, and more reliable clinical evidence—supporting innovation while strengthening patient protection and trust.

7. References

[1] International Council for Harmonisation (ICH). (2025). Guideline for Good Clinical Practice E6(R3) - Step 4 Final Guideline. ICH Secretariat. Retrieved from https://database.ich.org/sites/default/files/ICH_E6(R3)_Step4_FinalGuideline_2025_0106.pdf

[2] U.S. Food and Drug Administration (FDA). (2023). A Risk-Based Approach to Monitoring of Clinical Investigations—Questions and Answers (Guidance for Industry). FDA Center for Drug Evaluation and Research. Retrieved from https://www.fda.gov/regulatory-information/search-fda-guidance-documents/risk-based-approach-monitoring-clinical-investigations

[3] TransCelerate BioPharma Inc. (2013). A Rationale and Framework for Assessing the Risks and Benefits of Pharmacovigilance. Risk-Based Monitoring white paper. Retrieved from https://www.transceleratebiopharmainc.com/initiatives/risk-based-monitoring/

[4] Abhijit Bhatt. (2023). “The Revamped Good Clinical Practice E6(R3) Guideline: Major Amendments and Impact on Clinical Trial Conduct.” Perspectives in Clinical Research, 14(3), 215-221. PMC10679570. Retrieved from https://pmc.ncbi.nlm.nih.gov/articles/PMC10679570/

[5] Yuki Fuse, et al. (2025). “Artificial Intelligence in Clinical Data Analysis: A Review of Digital Twins, Synthetic Data, and Machine Learning Applications in Clinical Trials.” Data Science in Medicine, 15(4), 234-256.

[6] Hossein Akbarialiabad, et al. (2025). “Enhancing Randomized Clinical Trials with Digital Twins: Safety, Consent, Equity, and Real-Time Monitoring.” Nature Computational Science, March 21, 2024. PMC12494699. Retrieved from https://pmc.ncbi.nlm.nih.gov/articles/PMC12494699/

[7] Cyntegrity. (2023). FDA Releases Updated Risk-Based Guidelines for Clinical Trial Monitoring. Regulatory Intelligence Blog. Retrieved from https://cyntegrity.com/fda-releases-updated-risk-based-guidelines/

[8] http://cyntegrity.com/wp-content/uploads/2025/12/RBQM-Impact-Case-Study_by-Cyntegrity_DEC25.pdf?utm_medium=email&_hsenc=p2ANqtz-8NR8p9ycxRJjumtWf7f-Whh66bJloxFwUJ_3js52zOwlQzl1Ds3CozPIEXlH39sKUiBOeig7V0yR3WPbxW_xnSypq2koJxMvmhJn957fxAQpm4WRE&_hsmi=123581847&utm_content=123581847&utm_source=hs_automation

[9] U.S. Food and Drug Administration (FDA). (2021). Good Machine Learning Practice for Medical Device Development: Guiding Principles. FDA.

[10] European Medicines Agency (EMA). (2023). Artificial intelligence in the medicines regulatory network: roadmap and reflection materials. EMA.

Abbreviations

AE – Adverse Event	GCP – Good Clinical Practice
AI – Artificial Intelligence	GDPR – General Data Protection Regulation
API – Application Programming Interface	GMLP – Good Machine Learning Practice
CAPA – Corrective and Preventive Action	GxP – Good Practice (regulated quality standards)
CAD – Contextual Anomaly Detection	ICH – International Council for Harmonisation
CCPA – California Consumer Privacy Act	IQD – Intelligent Query Detection
CoU – Context of Use	IT – Information Technology
CRF – Case Report Form	KRI – Key Risk Indicator
CRA – Clinical Research Associate	ML – Machine Learning
CSA – Computer Software Assurance	NLP – Natural Language Processing
CtQ – Critical-to-Quality	PFS – Progression-Free Survival
CTMS – Clinical Trial Management System	QbD – Quality by Design
DMP – Data Management Plan	QMS – Quality Management System
DQ – Data Quality	QTL – Quality Tolerance Limit
DSS – Data Shared Service	RBAC – Role-Based Access Control
eCOA – Electronic Clinical Outcome Assessment	RBQM – Risk-Based Quality Management
EDC – Electronic Data Capture	RECIST – Response Evaluation Criteria in Solid Tumors
ePRO – Electronic Patient-Reported Outcome	RWE – Real-World Evidence
eSource – Electronic Source Data	SAE – Serious Adverse Event
EMA – European Medicines Agency	SDV – Source Data Verification
FTPS – File Transfer Protocol Secure	sFTP – Secure File Transfer Protocol
	SOP – Standard Operating Procedure

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Author Name: Priti Gupta
Company: GSK
Address: Bangalore, India
Work Phone: +91 9987111643
Email: priti.8.gupta@gsk.com
Website: https://www.gsk.com

Co Author Name: Dr. Deepa Balaji
Company: GSK
Address: Bangalore, India
Work Phone: +91 9886168019
Email: deepa.x.balaji@gsk.com
Website: https://www.gsk.com