ALGORITHMIC QUALITY ASSURANCE (AQA) AUDIT REPORT

Auditee: [REDACTED LIFE SCIENCES CORP.]

Facility Profile: Clinical Research & Diagnostic Technology Organization (Approx. 35 FTEs)

Audit Doctrine & Governance Standards: Algorithmic Audit Doctrine (AAD) 4-Pillar Governance Framework, ISO/IEC 42001 (Artificial Intelligence Management Systems), ISO/IEC 17025 (Testing and Calibration Competencies), FDA 21 CFR Part 11 / § 211.22 / § 211.194, EU GDPR Article 17, and Delaware Corporate *Caremark* Fiduciary Oversight Standards.

Executive Summary & Governance Scope

An independent Algorithmic Quality Assurance (AQA) readiness audit of [REDACTED] was conducted to evaluate the organizational deployment, runtime controls, and regulatory defensibility of enterprise AI tools and automated clinical data-processing pipelines. The audit evaluated operational risk across the Four Pillars of the Algorithmic Audit Doctrine (AAD): (1) Algorithmic Chain of Custody & GxP Isolation, (2) Algorithmic Competency & Blind Control Testing, (3) Human-in-the-Loop (HITL) Verification & Regulatory Defensibility, and (4) Cross-Functional AQA Governance & Director Oversight. The evaluation reviewed tenant permissioning configurations, LIMS database query interfaces, algorithmic change control logs, runtime audit trails, and executive governance workflows. While [REDACTED] demonstrates an innovative technical infrastructure and strong baseline compliance, substantial systemic remediation is required to bridge the gap between IT security assumptions and FDA Quality Control Unit (QCU) non-delegable oversight mandates.

System-by-System Evaluation Across the 4 Pillars

Pillar 1: Algorithmic Chain of Custody & GxP Data Isolation

Enterprise AI interfaces (including commercial large language models and Microsoft 365 Copilot integrations) were audited against internal GxP drives, structured laboratory databases (StarLIMS), and SharePoint document repositories. While commercial vendor SLAs guarantee tenant data isolation at the boundary, internal data permissioning structures revealed significant vulnerabilities regarding broad data indexing, lack of prompt-response cryptographic hashing, and insufficient technical boundaries separating unblinded clinical trial cohorts from automated data scrapers.

Pillar 2: Algorithmic Competency & Blind Control Testing

Analytical method automation, mass spectrometry peak-integration models, and data-screening scripts were evaluated against ISO/IEC 17025 competency standards. The organization utilizes standard vendor-supplied installation qualification (IQ/OQ) protocols; however, there is a total absence of continuous blind-control challenge testing, automated drift monitoring, and baseline regression benchmarking to evaluate model performance against shifting analytical data profiles.

Pillar 3: Human-in-the-Loop (HITL) Verification & Regulatory Defensibility

The site's operational workflows were audited against 21 CFR § 211.22(c) (non-delegable QCU review authority) and 21 CFR § 211.194 (complete data audit trails). Workflows utilizing automated AI summaries for batch record deviation assessments rely heavily on passive user concurrence rather than active, documented, and technically verifiable human-in-the-loop validation, exposing corporate officers to regulatory citations and *Caremark* duty-of-oversight liabilities.

Pillar 4: Cross-Functional AQA Governance & Director Oversight

Governance review assessed the communication channels and charter agreements between IT Architecture, Wet-Lab Quality Assurance, and Legal/Regulatory Compliance. Current AI initiatives are driven almost exclusively through IT procurement and software engineering silos, with minimal structural involvement from clinical QA or legal counsel, resulting in an Algorithmic Maturity Model standing at Level 1/2 (Ad-Hoc / Reactive).

Audit Findings & High-Risk Opportunities for Improvement

1. Major Finding (Internal Remediation Required; Non-Reportable)

- AAD Pillar: Pillar 1 (Algorithmic Chain of Custody & GxP Isolation) & Pillar 3 (HITL / 21 CFR § 211.22)
- Condition: During the audit of enterprise AI connections to internal network drives, it was discovered that an enterprise copilot tool had been granted read-level indexing access to a centralized SharePoint directory containing both unblinded Phase II clinical trial pharmacokinetic (PK) datasets and restricted human resource files. Furthermore, automated natural language queries executed by non-blinded biostatistical staff generated summary outputs that were logged into an unencrypted, ephemeral server cache without generating a time-stamped, permanent audit trail under 21 CFR § 211.194.
- Impact & Classification: Classified as a Major Internal Finding due to failure of access permissioning controls and non-contemporaneous data custody tracking. A retrospective security log audit confirmed that no unblinded data was exposed to unauthorized clinical site monitors, blinded investigators, or external third parties, and study blinding remained completely intact. As there was no breach of patient confidentiality or trial validity, this issue does not constitute a reportable violation to regulatory bodies or sponsors. However, immediate technical de-permissioning, tenant directory re-segmentation, and implementation of an Algorithmic Chain of Custody protocol are required prior to resuming automated data queries.

2. High-Risk Opportunity for Improvement (OFI 01: Model Drift & Blind Control Testing)

- AAD Pillar: Pillar 2 (Algorithmic Competency Assurance — ISO/IEC 17025)
- Observation: Machine learning models utilized for automated curve-fitting and chromatographic anomaly detection are not subjected to periodic blind-control challenge sets. It is strongly recommended to establish a formalized SOP mandating quarterly insertion of known, blinded historical data sets to quantitatively measure and document model drift, variance thresholds, and precision boundaries.

3. High-Risk Opportunity for Improvement (OFI 02: Passive Human-in-the-Loop Failure Modes)

- AAD Pillar: Pillar 3 (Human-in-the-Loop Defensibility — 21 CFR § 211.22(c))
- Observation: Out-of-specification (OOS) root-cause screening summaries generated by AI assistants are routinely approved via a single-click confirmation button in the electronic document system. To establish legally defensible human oversight and mitigate *Caremark* director liability, the SOP should require active, structured verification prompts where the reviewer must independently confirm critical data inputs and override parameters before sign-off.

4. High-Risk Opportunity for Improvement (OFI 03: StarLIMS Dynamic Query Injection Boundaries)

- AAD Pillar: Pillar 1 (GxP Data Isolation & Architectural Security)
- Observation: The API bridging natural-language AI queries to the production StarLIMS database lacks parameter sanitization and strict schema isolation, presenting a high risk of indirect prompt injection or unintended data joins. Implementing an isolated read-only intermediate data warehouse with strict data-masking layers is recommended to isolate primary GxP databases.

5. High-Risk Opportunity for Improvement (OFI 04: GDPR Right to Erasure vs. GxP Retention Conflict)

- AAD Pillar: Pillar 1 & Pillar 3 (Regulatory & Data Privacy Defensibility)
- Observation: The organization lacks a documented procedural reconciliation addressing the conflict between EU GDPR Article 17 (Right to Erasure) and FDA 21 CFR § 211.194 (Permanent Audit Trail Retention) regarding user prompts stored in model training logs. A formal Data Governance Policy must be drafted establishing cryptographic tokenization of user prompt histories to satisfy privacy rules without compromising GxP data integrity.

6. High-Risk Opportunity for Improvement (OFI 05: Algorithmic Change Control & Vendor Re-Training)

- AAD Pillar: Pillar 2 & Pillar 4 (Continuous Governance & ISO 42001 §8.4)
- Observation: Cloud-hosted third-party AI models utilized in document summarization receive automatic upstream foundation updates from vendors without internal change control evaluation. An Algorithmic Change Control Protocol should be instituted to quarantine model updates and execute automated regression testing suites before deploying updated model iterations into regulated workflows.

7. High-Risk Opportunity for Improvement (OFI 06: Chartering the Cross-Functional AQA Sub-Unit)

- AAD Pillar: Pillar 4 (Cross-Functional Governance & Boardroom Reporting)
- Observation: AI implementation and risk management currently reside solely within the IT infrastructure group, lacking formal oversight from Wet-Lab Quality Assurance and Legal Compliance. It is recommended to charter a formal, cross-functional Algorithmic Quality Assurance (AQA) Sub-Unit with direct quarterly reporting lines to the Board of Directors and Executive Leadership to advance the organization to Level 4 on the Algorithmic Maturity Model.

Conclusion & Strategic Roadmap

[REDACTED] possesses the advanced technical infrastructure required to establish an industry-leading, compliant AI operating model; however, bridging the gap between computational velocity and GxP regulatory defensibility requires immediate, structured cross-departmental coordination. To achieve Level 4 on the Algorithmic Maturity Model and eliminate executive *Caremark* exposure, executive leadership will immediately charter an Algorithmic Quality Assurance (AQA) Working Taskforce composed of designated leads from each operational discipline.

- Working Taskforce Assembly: Within 10 business days, leadership from IT & Data Engineering, Wet-Lab Quality Assurance (QCU), Legal & Regulatory Affairs, and Clinical/Analytical Operations will convene to formally establish the inter-departmental AQA steering charter.

- Phase 1 Coordination (Days 1–30 | IT & QA Lead): IT Architecture and Data Engineering will execute immediate de-permissioning of sensitive SharePoint GxP directories and implement parameter sanitization across StarLIMS API query pipelines. Concurrently, QA will verify complete prompt-logging and cryptographic audit trail capture compliant with 21 CFR § 211.194.
- Phase 2 Coordination (Days 31–60 | Clinical Ops & Legal Lead): Analytical Operations and Quality Assurance will author and implement Standard Operating Procedures for ISO 17025 algorithmic competency, establishing blind-control testing schedules for peak-integration models. Legal and Regulatory Affairs will draft the joint GDPR Article 17 / FDA Part 11 Data Governance Policy and formalize the active Human-in-the-Loop (HITL) review protocols.
- Phase 3 Coordination (Days 61–90 | Executive Committee Oversight): The cross-functional AQA Sub-Unit will be formally chartered into standard operational governance, establishing recurring bi-weekly algorithmic risk reviews and quarterly reporting cadence directly to the Board of Directors.

Formal Corrective and Preventive Action (CAPA) plans addressing the Major Finding must be submitted to the Quality Assurance Unit within 30 calendar days of this report's release. Continuous oversight, inter-departmental sign-offs, and verification of closed-tenant isolation boundaries must be finalized before authorizing expanded AI deployments across clinical study pipelines.