AI eSource Builder

Overview, Capabilities & Security Posture

Executive Summary:

The clinical trials industry is operating under increasing pressure: protocols are growing in complexity, study timelines are tightening, and the demand for data quality and regulatory traceability has never been greater. Against this backdrop, CRIO has developed the AI eSource Builder, a purpose-built, AI-powered platform module designed to automate and accelerate the translation of clinical trial protocols into structured study configurations. This is the first offering of the broader CRIO Innovation Studio.

This document provides a comprehensive overview of the CRIO AI eSource Builder, including what it does, how it works, and how it has been designed to meet the security, compliance, and governance expectations of sponsors, CROs, and their quality and technology oversight teams. It is intended for organizations evaluating the AI eSource Builder as part of a broader assessment of CRIO's platform or as a standalone capability review.

The AI eSource Builder operates exclusively on study design data, no patient data, no Protected Health Information (PHI) and is built around a non-bypassable human review gate before any AI-generated output reaches a production environment. It is governed by CRIO's established quality management system, its dedicated AI governance policy (POL-009), and a rigorous information security program that spans the full technology and vendor lifecycle.

What is the CRIO Innovation Studio?

The CRIO Innovation Studio is a toolbox CRIO has built for clients to better utilize the CRIO platform. It delivers the benefits of agentic AI in a compliant, controlled manner. The Innovation Studio houses a number of separate applications; the first and currently available application is the AI eSource Template Builder (also referred to as the AI eSource Builder).

The AI eSource Builder exists as a separate codebase from the core CRIO platform.

The separate codebase reflects the pace of AI development: CRIO needs the ability to iterate quickly as models evolve and as the pipeline is refined over time. Several software companies operating in the AI space have adopted a similar structure, maintaining an AI agentic suite that sits alongside, rather than within, the core product. This separation also serves a deliberate risk management purpose. By keeping AI workflows architecturally distinct from the core CRIO platform, any output generated by the AI eSource Builder remains confined to a separate environment until a human user makes an explicit decision to transmit it. Because AI is inherently probabilistic, this boundary ensures that the probabilistic outputs of the AI pipeline never interact directly with the core application. The human transmission step is not a formality; it is the architectural mechanism that keeps the AI eSource Builder's draft outputs from becoming consequential records until a qualified person has reviewed and approved them."

The Problem it Solves:

Setting up a study requires careful translation of protocol requirements into system configurations: visits, visit windows, procedures, eligibility criteria, and Schedule of Activities (SOA) mappings. This work is time-consuming, prone to transcription error, and must be repeated for each new study or protocol amendment. The SOA alone requires flipping between protocol pages and cross-checking footnotes to capture every visit correctly. This is a task that is both detail-intensive and poorly suited to manual effort at scale.
The AI eSource Builder performs this first pass automatically. Instead of starting from a blank template, study design staff begin from a draft that is already directionally correct. One that gives newer coordinators the same starting point as the most experienced staff. User time is then spent on clinical review and judgment, not transcription.

Core Functionality
The AI eSource Builder performs the following functions:

Capability	Description
Protocol & eCRF Ingestion	Accepts sponsor protocol documents and eCRF Guidelines as inputs. No PHI or patient-level data is involved at any stage.
Structured Extraction	Uses AI to extract arms, visits, visit windows, procedures, eligibility criteria, and SOA data, mapping them to eSource build configuration fields.
Build Rules Engine	Applies a library of pre-defined, deterministic build rules (Always, Keyword Match, SOA Match) to ensure consistency and reduce variability in AI-generated outputs. Rules can reference an organization's custom procedure library, ensuring builds reflect how the site or sponsor already works.
Draft Build Generation	Produces a structured draft eSource template ("AI Draft Build") for human review across multiple review steps.
Human Review & Publication	Users review the draft step by step, making edits before publishing. Only upon explicit user action does the build publish to the Publisher site for final configuration. The AI never publishes independently.

Key Design Principle
The AI eSource Builder is designed as a human-in-the-loop system. No AI-generated output reaches a production environment without explicit review and user-initiated publication. This is

not a configurable setting. It is a non-bypassable architectural requirement.

How It Works: A Closer Look at the Technology

Foundation Models

The AI eSource Builder uses Claude (Anthropic) and OpenAI APIs as its underlying large language model (LLM) infrastructure. CRIO has engaged both providers under contractual no-training agreements, which explicitly prohibit the use of any customer data (including protocol documents, study configurations, or extracted data) for model training or fine-tuning. These providers are subject to CRIO's vendor management program and periodic vendor review.

Structured Outputs and Deterministic Logic

CRIO does not rely solely on LLM outputs for structured decisions. The AI eSource Builder employs two important safeguards against the inherent variability of generative AI:

- Pydantic structured output schemas: LLM outputs are constrained to validated formats through Pydantic models, preventing free-form or malformed responses from entering the build pipeline.
- Build rules engine: A deterministic rules engine applies explicit logic where appropriate, categorizing procedures using Always, Keyword Match, and SOA Match rules, reducing reliance on LLM inference for decisions that can be made deterministically. These rules can be configured to reflect an organization's own custom procedure library.

Accuracy and Realistic Expectations

AI is inherently probabilistic: it makes inferences from a wide variety of inputs (e.g. protocols, eCRF guidelines) that can vary significantly in format and terminology. The goal of the AI eSource Builder is to be directionally correct and to save meaningful time, not to produce a build requiring no review. The more the tool is used, the more CRIO can refine and improve the pipeline over time.

This framing is important for evaluators to understand:

The tool is not presented as a replacement for clinical judgment. It is presented as a productivity multiplier for experienced study design staff who retain full review authority over every build.

Explainability and Source Traceability

Each AI-generated element in the draft build is traceable back to its specific location in the source protocol document. Reviewers can verify every extracted item against the original text using the document reviewer functionality that allows for a side-by-side comparison of the source document and draft build. The review interface surfaces the build rules that drove each procedure decision, and an admin-accessible “Builder Logic” page exposes the underlying LLM prompts used at each extraction step, providing full transparency into how the model interpreted the protocol.

Handling Ambiguity and Uncertainty

Where the AI cannot confidently match a Schedule of Activities item, it explicitly surfaces these as “Unmatched SOA Placeholders” for human resolution rather than making a best-guess assignment. This ensures that ambiguity is surfaced to the reviewer rather than silently propagated into the build.

Audit Trails and Observability

The AI eSource Builder maintains comprehensive audit logging at multiple levels:

- Every change made during the review process is captured with user attribution and timestamp.
- The system distinguishes between “AI Draft Build” and “Final Build” states, preserving a clear record of what the AI generated versus what the human reviewer approved.
- LLM call observability is captured via OpenTelemetry, Langfuse, and Logfire for latency tracking, output inspection, and model performance monitoring.

Data Scope: What the AI eSource Builder Does and Does Not Process

One of the most important considerations for any AI system operating in the clinical trials environment is the nature of the data it processes. The AI eSource Builder has been deliberately scoped to ***exclude*** patient-level data and PHI from its pipeline.

Data Category	In Scope / Out of Scope
Sponsor protocol documents	IN SCOPE, the primary input to the system.
eCRF Guidelines	IN SCOPE, accepted as a supplementary input to the Source Template Builder.
Study design configuration data	IN SCOPE, arms, visits, visit windows, procedures, SOA mappings, eligibility criteria.
Protected Health Information (PHI)	OUT OF SCOPE, PHI never enters the AI pipeline at any stage.
Patient-level data	OUT OF SCOPE, no patient records are processed or stored by the AI eSource Builder.
Identified site or subject data	OUT OF SCOPE, the AI eSource Builder operates exclusively on protocol and study design artifacts.

This deliberate data scoping significantly limits the privacy, confidentiality, and regulatory exposure associated with the AI eSource Builder's use, and simplifies the assessment of applicable data protection obligations for client organizations.

Security Architecture and Controls

The AI eSource Builder operates within CRIO's broader information security program, governed by POL-002 Security Policy and managed through an established set of standard operating procedures covering identity management, infrastructure, incident response, vendor management, and business continuity.

Infrastructure

The AI eSource Builder runs entirely on Google Cloud Platform (GCP) infrastructure. GCP provides foundational security controls including AES-256 encryption at rest, TLS 1.2 or higher for data in transit, and encryption key management via Google Keystore with per-operation access logging. GCP undergoes regular independent third-party audits against ISO 27001, ISO 27017, ISO 27018, SOC 2, NIST 800-53, HIPAA, and other relevant frameworks.

Access Control

Access to the AI eSource Builder is not automatically granted to all CRIO platform users. Client organizations must explicitly request user access, which is then provisioned on a named-user basis. CRIO applies Least Privilege and Need-to-Know principles to all system access. Any

organizations with multi-factor authentication enabled inside the CRIO platform will also inherit the same security and access configuration for the AI eSource Builder.

Vendor Security

CRIO's vendor management program (**SOP-QMS-005 Vendor Management**) governs all third-party relationships, including AI foundation model providers. Anthropic and OpenAI are managed as vendors subject to risk classification, vendor assessments, and periodic review. No-training agreements are contractually in place with both providers. Google's compliance posture is reviewed annually.

Vulnerability Management and Monitoring

CRIO's vulnerability management program (**SOP-IT-006**) includes SIEM/SOAR monitoring via Google Security Command Center Enterprise, continuous log review, periodic penetration testing by an independent third-party firm at minimum annually, and security scanning against OWASP 2017 and 2021 standards. Engineering staff receive periodic secure coding training based on the OWASP Top 10.

AI-Specific Risk Controls

Risk	Mitigation
Prompt injection	Structured output schemas (Pydantic) constrain LLM responses to validated formats; build rules engine applies deterministic logic independent of LLM outputs.
Data poisoning / model inversion	No PHI or patient data enters the pipeline; no-training agreements with Anthropic and OpenAI prohibit use of customer data for model training.
Hallucination / inaccurate extraction	Non-bypassable human review gate; source traceability to protocol document; explicit surfacing of unmatched items; accuracy benchmarks against human ground truth; golden test fixtures to catch regressions.
Model drift	LLM call observability via OpenTelemetry, Langfuse, and Logfire; accuracy benchmarks tracked over time; evaluation/scoring system for ongoing build quality measurement on roadmap.
Unauthorized access	Role-based access control; MFA enforcement; access provisioned only on explicit client request; least-privilege principles.

AI Governance

CRIO maintains a dedicated AI governance policy (**POL-009 Artificial Intelligence at CRIO**) that applies to all staff, contractors, and anyone using AI on behalf of CRIO. **POL-009** covers responsible and ethical use, bias identification and mitigation, data protection obligations specific to AI use, transparency and accountability standards, training requirements, explicit prohibitions on misuse, and the requirement that all AI-generated content be reviewed and approved by a human before use in decision-making.

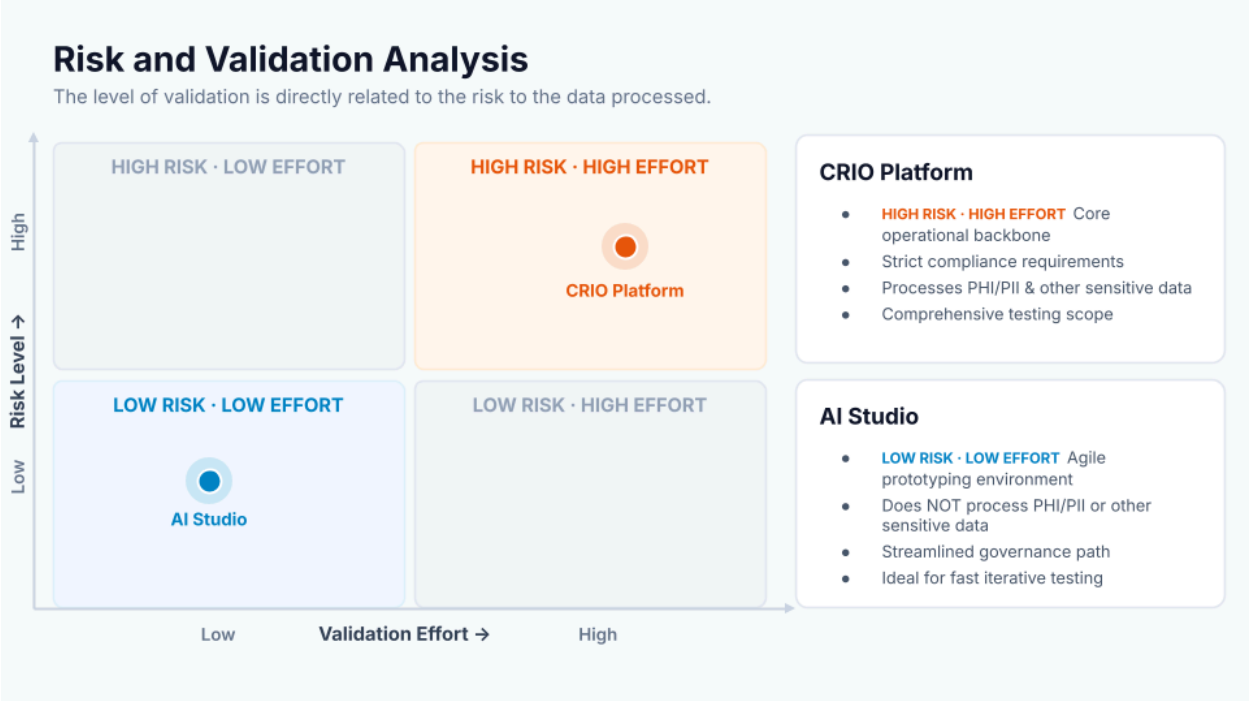
Alignment with Emerging Regulatory Frameworks
The joint FDA-EMA Guiding Principles of Good AI Practice in Drug Development were published in January 2026. CRIO's human-in-the-loop design, audit trail architecture, and AI governance policy are aligned with the core principles these frameworks establish, including human-centric design, risk-based approaches, data governance, and transparent accountability for AI-assisted processes.

Quality Management and Change Control

The AI eSource Builder is subject to CRIO's established Quality Management System (QMS), which governs the full lifecycle of the platform including development, release, incident management, and vendor oversight.

Software Development and Release

The software development process and validation activities leverage some of the best practices and tools available for the CRIO platform.
It is important to note that the AI eSource Builder is firmly placed *outside* of the CRIO platform and *no GxP data exists* within the AI eSource Builder, the level of validation and CRIO's validation approach as a whole differs from the CRIO platform.
Because of the differences in risk identified (in alignment with frameworks such as ICH E6(R3) Good Clinical Practices), the level of validation and verification performed for each release differs from what clients have seen for the CRIO Platform.



What Checks Are Performed After An Update

The following checks run automatically on every code change (every pull request and every merge to main):

- About 3,100 automated tests (unit and integration). They cover the build logic, the CRIO API payload, organization access control and the database models.
- Linting and type checking on the whole codebase.
- A database migration check. It blocks any deploy whose schema changes would conflict.

Promoting a change to the Production environment is done manually, requiring a human-in-the-loop approach - this is intentional to ensure a human has the final authority to push something to Production.

The automated tests run on every code update. Beyond that, CRIO follows a risk-based approach. Any ticket that may impact one of three areas gets a "**verify:critical**" tag:

- The CRIO integration, or anything moving data between systems
- Data pipelines and transformation logic
- Authentication and access control

Those tickets have a test plan outlined inside the ticket itself and a human review performed.

Tools and Artifacts used during the development and verification of AI eSource Builder

- **GitHub Actions**: test, lint and type check results for every change
- **GitHub pull requests**: the review record for every change
- **Jira**: Where ***verify:critical*** tickets with their test plans and evidence are created and stored
- **Langfuse**: prompt version history and the evaluation results for each prompt change
- **The app's own audit log**: records who changed what in a build

AI Pipeline Change Control

The AI eSource Builder uses golden test fixtures, a curated set of benchmark protocols with known ground-truth outputs to validate extraction accuracy before any pipeline change ships. Any fix must generalize across all protocols in the test suite, not merely the specific case that prompted the change. This prevents targeted fixes that inadvertently degrade performance on other protocol types or therapeutic areas.
CRIO's Study Design team serves as an active user of the AI eSource Builder and a significant internal stakeholder shaping the product's development. Improvements are made on an ongoing basis and informed by user feedback as well as internal usage experience.

Incident Management

Incidents involving the AI eSource Builder are handled through CRIO's incident management process (**SOP-QMS-003 Incident Management and Reporting**) and quality event management process (**SOP-QMS-002 Quality Event Management**), including formal CAPA where required. **POL-009 Artificial Intelligence at CRIO** requires all staff to report suspected AI misuse or output quality concerns to their supervisor or the Data Protection Officer.

The Human-in-the-Loop Imperative

In a highly regulated environment, the introduction of AI into any workflow that touches clinical trial data demands a clear and enforceable boundary between AI-generated outputs and human-verified conclusions. The AI eSource Builder was designed from the outset with this principle as a non-negotiable architectural requirement.
Trust in the AI eSource Builder does not mean the AI gets it right and the user signs off blind. It means the tool is built so the user stays in control. The AI works within a library of pre-defined build rules, so it is not free-associating off protocol text. It is applying the same structured logic every time. Those rules can reflect an organization's own procedures. The question is not whether to trust the AI outright; it is whether the user trusts themselves to catch what needs

adjusting, and the tool is designed to make that review fast rather than requiring the output to be taken on faith.

The Review Workflow

Stage	Description
1. AI Draft Build	The AI processes the protocol and eCRF Guidelines and generates a structured draft build. This output is clearly marked as an AI Draft Build and is not accessible to production environments.
2. Step-by-Step Human Review	The user reviews the draft across multiple defined review steps, verifying extracted elements against the protocol, editing where necessary, and resolving any Unmatched SOA Placeholders.
3. User-Initiated Publication	Only upon explicit user action does the build publish to the Publisher site. The AI never publishes independently.
4. Publisher Configuration	Final configuration is completed by the user in the Publisher environment before publication to the site(s) that will be collecting data. The AI eSource Builder output is a starting point, not a final product.

This design ensures that the AI eSource Builder functions as a productivity tool for experienced study design staff, not as an autonomous decision-maker. Clinical expertise remains the controlling factor in every build.

Validation: A Risk-Based Assessment

Validation is among the first questions any CRIO client will raise when assessing a new software tool in a regulated environment. The short answer for the AI eSource Builder is that CRIO applies a risk-based, proportionate approach to validation. This method is grounded in both established GCP guidance and CRIO's own documented risk-management procedures. The AI eSource Builder's specific design characteristics position it as a low-risk tool under that framework.

Proportionate Controls

The governing principle is articulated in ICH E6(R3) §3.10.1.3 (Risk Control): risk control should be proportionate to the importance of the risk to participants' rights, safety and well-being and to the reliability of trial results. This is the guideline's explicit authorization to scale validation controls to the actual risk profile of a given tool or function.

The criteria for evaluating that risk profile are set out in ICH E6(R3) §3.10.1.2: the likelihood of harm, its detectability, and its impact on participant protection and result reliability. These three criteria map directly to CRIO's own risk-management procedure (**SOP-DEV-002 Risk Management**), which defines Risk Priority as Risk Class (Severity × Probability) × Detectability. CRIO's internal procedure has been developed to align directly with the expectations of the ICH risk-management principle.
The FDA's Computer Software Assurance guidance (finalized September 2025) reinforces this position by endorsing a risk-based approach to software validation that moves away from exhaustive documentation toward critical thinking and fit-for-purpose evidence. Validation stringency should be commensurate with the risk and context of use of the system.

Why the AI eSource Builder is a Low-Risk Tool

Evaluated against the ICH E6(R3) three-part criteria, the AI eSource Builder presents a low residual risk profile for two structural reasons that are built into its architecture.

Risk Factor	AI eSource Builder Position
No data entry into regulated records	The AI eSource Builder does not write to any clinical trial record, eSource, or CTMS in production. It generates a draft study build configuration that exists in a separate environment until explicitly published by a human user. No trial data, patient data, or source document is created or modified by the AI.
Non-bypassable human review gate	Every AI-generated output must pass through an explicit, multi-step human review process before it can be published. The AI never acts on a production environment autonomously. This means a human with clinical expertise is independently re-evaluating every element before it becomes consequential. Any error introduced by the AI is therefore highly detectable and correctable before it affects the reliability of trial results.
Low model influence	The AI is one input among several: it is constrained by the build rules engine, the user's clinical judgment, the organization's custom procedure library, and the final Publisher configuration step. The AI's contribution to the final study build is mediated, not determinative.
No impact on participant safety at point of AI action	The AI operates exclusively on protocol documents and study design artifacts. It has no interaction with participant-level data, consent records, or safety data at any point in its workflow.

Applied to the ICH 3.10.1.2 criteria: the likelihood of harm from an AI extraction error is low because no regulated record is created until a human approves it; the detectability of any error is high precisely because the human review gate independently re-adjudicates the output; and

the impact on result reliability is low because the authoritative record, the eSource entry, the study configuration originates from the human-reviewed Final Build, not from the AI Draft. This mirrors the FDA's own two-dimensional risk matrix for AI tools: ***Model Influence*** (the AI's contribution relative to other evidence) is low, because human review is the controlling factor; and ***Decision Consequence*** (the significance of an adverse outcome from an incorrect AI output) is low, because no adverse outcome can reach a production environment without passing through that human gate.

What This Means for Client Organizations

Clients evaluating the AI eSource Builder do not need to treat it as a high-risk, GxP-critical system requiring independent revalidation. Under the FDA's guidance on computer software assurance, clients should:

- Review CRIO's validation documentation and release risk assessments, which are available upon request from the CRIO Document Center.
- Confirm that CRIO's SOPs and change control processes are adequate for the intended use consistent with the approach endorsed in ICH E6(R3) and the FDA's guidance on leveraging vendor validation evidence.
- Document internally that the AI eSource Builder's risk profile was assessed and found proportionate to the controls in place, mirroring the rationale described in this section.
- Note that no patient data, no eSource records, and no regulated clinical trial data passes through the AI eSource Builder pipeline, which limits the scope of applicable predicate rule obligations.

Clients with more stringent internal validation requirements, or sponsors whose own quality frameworks impose additional obligations, are encouraged to contact CRIO's Compliance team to discuss the available documentation and any supplementary evidence that may support their internal assessment.

Access Provisioning

Client organizations can request access provisioning for named users. Access is not self-service and is not automatically granted. The access request process ensures that only personnel with a legitimate need are granted access to study build data within the AI eSource Builder.

Questions About Subprocessors

Anthropic and OpenAI serve as data processors for AI inference within the AI eSource Builder. Both operate under no-training agreements prohibiting use of customer data for model improvement. Because the AI eSource Builder does not process PHI, the scope of data

processor obligations is limited; however, all third-party AI providers are governed by CRIO's vendor management framework (**SOP-QMS-005**). Any engagement involving potential PHI exposure would require an executed Business Associate Agreement per **SOP-OPS-007**.

Conclusion

The CRIO AI eSource Builder represents a deliberate, governed approach to applying generative AI within the clinical trials environment. It was built to solve a real operational problem: the manual burden of translating complex protocols into structured CTMS configurations, while respecting the compliance and data integrity requirements of the industry it serves.

Its design reflects three core commitments:

- that AI should augment, not replace, clinical expertise;
- that every AI-generated output must be traceable, explainable, and subject to human approval before it reaches a production environment; and
- that the security and governance framework surrounding the tool must meet the same standards expected of the broader platform it integrates with.

For organizations evaluating the AI eSource Builder as part of a broader CRIO assessment, or as a standalone capability review, CRIO welcomes further dialogue and is available to support your due diligence process.

For additional information please contact:

Marc Wartenberger
VP, Compliance & Security
marc.wartenberger@clinicalresearch.io

Annex 1: Common Evaluator Questions

The following questions and answers reflect the most common inquiries received from security, compliance, and technology teams during AI eSource Builder evaluations. They are intended to supplement the sections above with direct, question-specific responses.

What type of AI models or algorithms are used (e.g., predictive analytics, natural language processing, computer vision)?
The AI eSource Builder uses large language models (LLMs), specifically the Claude API (Anthropic) and OpenAI API to perform natural language processing on clinical trial protocol documents. The system also incorporates a deterministic build rules engine that applies explicit rule-based logic where structured decisions can be made without LLM inference, and Pydantic structured output schemas that constrain LLM outputs to validated formats. The system does not use predictive analytics, computer vision, or diagnostic AI.

What data does the AI system use, and is it patient-related?
The AI eSource Builder processes sponsor protocol documents and eCRF Guidelines only. No patient data, no Protected Health Information (PHI), and no subject-level data of any kind enters the AI pipeline at any stage. The only data stored within the AI eSource Builder is study build configuration and sponsor protocol documents. This is a deliberate architectural decision that significantly limits the privacy and regulatory exposure associated with the system.

How is patient data protected, and does the system comply with regulations (e.g., GDPR, HIPAA)?
Because the AI eSource Builder does not process patient data or PHI, the direct exposure under patient data protection regulations is limited by design. All data in transit is encrypted using TLS 1.2 or higher; data at rest is encrypted with AES-256. Any third party with potential PHI exposure is required to execute a Business Associate Agreement.

Has the AI model been tested for bias, and how is fairness ensured?
The AI eSource Builder operates on protocol documents rather than patient-level data, which means traditional demographic bias auditing is not directly applicable. Quality assurance focuses on extraction accuracy across diverse therapeutic areas, study phases, and protocol formats to ensure consistent performance regardless of study type. Golden test fixtures validate that pipeline changes generalize across the full range of protocol types, not just the specific case that prompted a given change. CRIO's AI governance policy (**POL-009**) requires staff to evaluate AI outputs for bias and inaccuracies and mandates collaboration with providers to address identified issues. All AI-generated content must be reviewed and approved by a human before use.

Can clinicians review and override AI-generated decisions?
Yes, and this review is non-bypassable. The AI eSource Builder is explicitly designed as a human-in-the-loop system. The AI generates a draft build clearly marked as an "**AI Draft Build**" which is inaccessible to production environments. Users review every extracted element across multiple review steps, verify against the source protocol, edit as needed, and initiate publication explicitly. The AI never publishes independently. Clinical expertise remains the controlling factor in every build.

Does the system provide explanations for AI-driven recommendations?
Yes. Each AI-generated element is traceable to its specific location in the source protocol document, allowing reviewers to verify extraction against the original text. The review interface surfaces the build rules (Always, Keyword Match, SOA Match) that drove each procedure decision. An admin-accessible "Builder Logic" page exposes the underlying LLM prompts used at each extraction step. Where the system cannot confidently match a SOA item, it explicitly flags these as "Unmatched SOA Placeholders" for human resolution rather than making an assumption.

How is AI accuracy measured, and what are its error rates (false positives/negatives)?
AI extraction accuracy is benchmarked against human ground truth using golden test fixtures (curated protocols with known correct outputs). These serve as regression tests: any pipeline change must be validated against the full fixture set before release. Accuracy metrics are tracked over time. An expanded evaluation and scoring system for ongoing build quality measurement is on the near-term product roadmap. It is also important to note that the AI eSource Builder is designed to be directionally correct rather than perfect. The non-bypassable human review gate functions as an independent quality control layer that catches extraction gaps before they reach a production environment.

Has the AI been clinically validated for use in healthcare?
The AI eSource Builder is a study configuration tool, not a clinical decision support system. It does not provide diagnostic, prognostic, or treatment recommendations and is not subject to clinical validation requirements applicable to Software as a Medical Device (SaMD). CRIO applies a risk-based, proportionate approach to validation grounded in ICH E6(R3) §3.10.1.3 and the FDA's Computer Software Assurance guidance. The AI eSource Builder's architecture (no direct data entry into regulated records, a non-bypassable human review gate, and low model influence) supports a proportionate validation classification. A full discussion of the validation rationale, applicable controls, and what this means for client organizations is provided in the "**Validation: A Risk-Based Assessment**" section of this document.

How frequently is the AI model updated or retrained?

CRIO does not train or retrain the underlying foundation models. Claude (Anthropic) and OpenAI are engaged under no-training agreements prohibiting the use of customer data for model training or fine-tuning. Foundation model updates are managed by Anthropic and OpenAI respectively. CRIO manages its own AI pipeline, including prompt logic, build rules, and structured output schemas through its software development and change control process. Improvements are made on an ongoing basis, informed by user feedback and internal Study Design team experience.

How are AI-related risks (e.g., incorrect predictions) identified and managed?
AI-related risks are managed through multiple layered controls. At the design level: structured output schemas constrain LLM responses to validated formats; the deterministic build rules engine reduces reliance on LLM inference; and unmatched or ambiguous items are explicitly surfaced to reviewers. At the operational level: the non-bypassable human review gate ensures no AI output reaches production without approval; golden test fixtures catch regressions before release; and LLM call observability via OpenTelemetry, Langfuse, and Logfire enables ongoing monitoring. At the governance level: **POL-009** requires staff to report suspected misuse or quality concerns; incidents are handled through **SOP-QMS-003** and **SOP-QMS-002**, including formal CAPA where required.

Who is responsible for AI maintenance, updates, and monitoring?
Responsibility is distributed across defined CRIO roles. The Engineering team owns pipeline development, infrastructure, and release management. The TIP (Technology, Infrastructure & Platform) team manages infrastructure monitoring, vulnerability management, and security operations. The Study Design team serves as an active user and key internal stakeholder. Corporate QA & Compliance owns the governance framework, including **POL-009**, the QMS, and vendor oversight for Anthropic and OpenAI. The Data Protection Officer coordinates incident response for data-related events.

What service level agreements (SLAs) exist for AI performance and issue resolution?
SLA terms are defined in client agreements and may vary by engagement type. Prospective clients are encouraged to discuss specific SLA requirements during the contracting process. At a process level, CRIO's incident management framework (**SOP-QMS-003**) defines structured handling stages (Discovery, Triage, Investigation, Notification, Mitigation, Closure, and Effectiveness Monitoring) and client notification occurs within 72 hours of a confirmed data incident. Quality events implicating AI performance are subject to the formal CAPA process under **SOP-QMS-002**.

How does the vendor handle AI model drift and system degradation over time?
Model drift and degradation are monitored through several mechanisms. LLM call observability via OpenTelemetry, Langfuse, and Logfire provides ongoing visibility into latency, output quality, and anomalous behavior. Accuracy benchmarks are tracked over time against human ground

truth, and golden test fixtures are run against every pipeline change to detect regressions before they reach production. Because CRIO does not control the underlying foundation models, changes introduced by Anthropic or OpenAI to their APIs are monitored through observability tooling, and any degradation in extraction quality is flagged through the accuracy benchmarking process and addressed through CRIO's SDLC change control workflow.