

ML LABS INTELLIGENCE

Building a Clinical AI Report Signing Workflow

ML LABS engineered the complete report lifecycle for a healthcare AI platform — from AI-generated preliminary results through clinician review, signed confirmation, and automated EHR delivery.

CASE STUDY

Author Omar Trejo

Date 2026-01-05

ML LABS

mllabs.com/intel/ai-diagnostic-report-workflow

A healthcare AI platform needed more than accurate diagnostic results. It needed a complete clinical report workflow — the entire pipeline from AI-generated preliminary findings through clinician review, mandatory change tracking, digital signing with audit logs, PDF generation, and automated result delivery back to the hospital's EHR. Every step carries regulatory weight under 21 CFR Part 11, and a single gap in the audit trail invalidates the entire chain.

ML LABS designed and built this workflow end to end. The engagement covered the report lifecycle with enforced transitions, batch confirmation for high-volume clinician workflows, mandatory change reason tracking, PDF generation with clinical formatting requirements, and automated result delivery triggered on confirmation. The result is a system where every modification from AI preliminary to signed report is traceable, attributable, and defensible under audit.

The Problem

The platform's AI models produced accurate diagnostic results. But between a model output and a legally signed clinical report sits a dense layer of workflow engineering that most teams underestimate. [Research on implementing AI in clinical workflows](#) identifies the peri-implementation phase — the processes wrapping the model — as where most complexity concentrates.

The specific requirements were unforgiving. Every edit to an AI-generated report needed a mandatory change reason. Clinicians needed to review, confirm, and sign reports in batch without losing state. Signed reports had to lock permanently, with any unlock requiring a separate privileged action and its own audit trail. PDF output had to render correctly across variable report lengths with institution-specific branding. And confirmed reports had to trigger automated result delivery to the EHR — without sending duplicates when both preliminary and final delivery paths were active.

Report Lifecycle

ML LABS built a complete report lifecycle with full audit compliance. Every report passes through defined stages from AI-generated preliminary through clinician review, confirmation, signing, locking, and delivery — with the system enforcing valid transitions and preventing any sequence violation. This is not a status flag on a database row. It is a compliance-grade workflow where every boundary carries enforced preconditions and auditable side effects.

Mandatory Change Tracking

When a clinician modifies any AI-generated value — a measurement correction, an interpretation override, an artifact removal — the system requires a structured change reason before allowing confirmation. This satisfies **21 CFR Part 11 requirements for electronic records**: every modification is independently recorded with the original value, new value, clinician identity, reason, and timestamp — sufficient detail to reconstruct the complete history under audit.

Batch Confirmation for High-Volume Workflows

In production clinical environments, clinicians do not confirm reports one at a time. ML LABS built a confirm-and-next workflow that lets a clinician review, confirm, and advance to the next study in a single action — with every confirmation verified server-side before advancing. The batch workflow was hardened through production clinical use, catching and resolving edge cases that only surface under real clinical volume and timing conditions.

Signing and Locking

Signing is the legal act that seals the report. After signing, the report enters a locked state where no further edits are permitted. Any unlock requires a separate privileged action with its own mandatory reason and audit trail, creating a complete chain of custody from first AI output to final signed document.

// The difference between a clinical AI prototype and a production system is not the model. It is whether every change to the model's output is tracked, attributed, and defensible under audit.

ML LABS hardened the lock enforcement through iterative testing — catching and resolving cases where the locked state could be circumvented through specific navigation paths. These are the kinds of compliance-critical issues that only surface through exhaustive testing of real clinical workflows, not theoretical review.

Prior Report Navigation

Clinical AI interpretation does not happen in isolation. Clinicians compare current findings against prior studies for the same patient — longitudinal context that changes the clinical significance of individual measurements. ML LABS built a prior report navigation system that shows previous report titles and enables direct comparison, with the current and prior panels scrolling in lockstep and handling reports of different lengths gracefully.

PDF Generation

Clinical PDFs are not straightforward rendering jobs. They are regulated documents that must render identically across systems, embed all required assets, and handle variable content gracefully.

ML LABS built the PDF generation pipeline to handle variable report lengths without layout failures, institution-specific branding with configurable cover page layouts, consistent multi-page formatting with correct headers and footers, and self-contained output that works in network-restricted clinical environments. New institutions are supported through configuration rather than code changes.

Result Delivery

Confirmation triggers automated result delivery back to the EHR via HL7 result messages. The delivery timing is critical: sending results before the report is fully signed creates compliance exposure, but deferring too long delays clinical action.

ML LABS built result delivery with support for both preliminary and final result paths, with per-organization configuration controlling which paths are active. The system prevents duplicate sends across pathways and handles PDF delivery to EHR document management systems. The [companion article on HL7 messaging](#) covers the message construction layer in detail.

Boundary Condition

This workflow model assumes the AI system produces structured, discrete results that a clinician can meaningfully review field by field. If the AI output is an opaque score or free-text narrative without structured fields, the change-tracking and

confirmation flow described here does not apply directly. The lifecycle enforcement still matters, but the change reason taxonomy needs different design when review is qualitative.

Similarly, some screening applications deliver results without individual clinician sign-off. In those cases, the full lifecycle is unnecessary overhead. The right workflow complexity must match the regulatory and clinical requirements of the specific use case.

First Steps

1. **Map every report state and transition before writing code.** Define the preconditions, side effects, and audit requirements for each state change. The lifecycle design determines whether the system is auditable and recoverable — retrofitting it later is expensive and error-prone.
2. **Build the audit infrastructure before the review UI.** Mandatory change tracking and the validation layer that enforces compliant transitions are the foundation. Everything else depends on them.
3. **Test production edge cases explicitly.** Accidental confirmations, unconfirm-reconfirm flows, batch workflows under network latency, and comparison views are all testable. The measure of readiness is whether every edge case has been exercised and logged correctly.

Practical Solution Pattern


Design the clinical report workflow as a compliance-grade lifecycle with enforced transitions, mandatory change tracking, and controlled delivery timing that prevents duplicate result sends. Build the audit infrastructure first, enforce change reasons as

preconditions rather than optional metadata, and invest in the edge cases — batch confirmation, lock enforcement, comparison views — that only surface under real clinical use. PDF generation must handle variable content lengths, institution-specific branding, and multi-page layouts without manual intervention.

This approach works because the primary failure mode in clinical AI report systems is not inaccurate AI output — it is untraceable changes to that output. When every modification is logged with its reason, timestamp, and operator identity, the system satisfies regulatory requirements by design rather than by retrofit. Getting this right the first time matters more than in traditional software — clinical AI systems that launch with audit gaps erode organizational confidence in the entire initiative and trigger expensive remediation cycles. If the clinical workflow is defined and the constraint is production-quality engineering, **AI Workflow Integration** is the direct build path.

References

1. Wang, F., Beecy, A. **Implementing AI Models in Clinical Workflows: A Roadmap**. *BMJ Evidence Based Medicine*, 2025.
2. U.S. Food and Drug Administration. **Part 11, Electronic Records; Electronic Signatures — Scope and Application**. *Regulatory Reference*, 2003.
3. Wu, D. T. Y., et al. **Using EHR Audit Trail Logs to Analyze Clinical Workflow**. *AMIA Annual Symposium Proceedings*, 2018.
4. Monico, E., Schwartz, I. **Communication and Documentation of Preliminary and Final Radiology Reports**. *Journal of Healthcare Risk Management*, 2010.
5. U.S. Food and Drug Administration. **Software as a Medical Device (SaMD)**. *Regulatory Reference*, 2024.
6. HL7 International. **FHIR R5 Specification**. HL7 International.
7. Regenstrief Institute. **LOINC and Other Standards**. LOINC.



ML LABS

Custom AI Systems for High-Value Workflows

mllabs.com