The 510(k) Gap

FDA-cleared radiology AI is running in 1,100+ hospitals. The 510(k) clearance covers the algorithm at the time of the study. When the vendor updates the model, most hospitals have no protocol to detect it, validate it, or delay it.

01

Background

How FDA-cleared radiology AI became the most deployed — and least governed — clinical AI in enterprise health systems

01 / BACKGROUND

The Radiology AI Deployment Landscape

- FDA has cleared 950+ AI/ML-based Software as a Medical Device (SaMD) through Q1 2025 — approximately 75% are imaging-based radiology applications
- Viz.ai: 1,100+ hospitals, 20 FDA clearances — detects LVO stroke, pulmonary embolism, aortic dissection; pages on-call team directly without standard dictation queue
- Aidoc: 1,200+ health systems — continuous AI triage across CT, X-ray, MRI for 20+ pathologies including incidental findings
- Nuance PowerScribe 360 with AI: 10,000+ radiology practices — AI-assisted finding detection embedded in radiologist dictation workflow
- Deployment driver: U.S. radiologist shortage. ~35,000 practicing radiologists; imaging volume up 40% in a decade. AI deployed to close the gap — governance has not kept pace.

The FDA 510(k) Model Gap

- 510(k) clearance covers: the algorithm version validated in the study, the imaging modalities and scanner configurations, and the patient population and clinical setting at time of clearance
- Predetermined Change Control Plan (PCCP): FDA's framework for managing post-clearance algorithm updates — voluntary, not consistently required or submitted by vendors
- When vendor updates the model: hospital continues operating under original clearance while running a materially different algorithm; no mandatory customer notification requirement under current FDA rules
- Most hospital radiology AI contracts omit: (a) advance model update notice, (b) version-specific performance data for local scanner fleet, (c) contractual right to delay an update, (d) model version change log access
- Result: most health systems cannot confirm what algorithm version is running in production today or whether it matches their FDA clearance documentation

01 / BACKGROUND

Alert Fatigue and Automation Bias

- Studies at two academic medical centers: deploying radiology AI without threshold calibration increased radiologist alert volume 35–55% in the first 90 days
- Within 6 months: radiology staff reported reviewing alerts at reduced attention levels — the behavioral shortcut that defines alert fatigue
- IRB-approved automation bias studies: radiologist independent detection rate for flagged finding classes declined 12–18% when AI alerts were present in the worklist interface
- The AI was deployed to improve detection accuracy. In conditions of alert fatigue and automation bias, the deployment degraded it.
- Liability framework: radiology AI is classified as clinical decision support — the radiologist who signs the report bears liability for missed findings regardless of what the AI flagged

02

Decision Required

Your health system has deployed FDA-cleared AI to your radiology department. Your procurement team validated a clearance number. Three governance questions remain unresolved.

First: what algorithm version is running in production right now — and does it match the clearance documentation your compliance team filed? Most health systems cannot answer this question without contacting the vendor.

Second: who is liable when the AI misses a finding and the radiologist signs the report? Your professional liability insurer may not know you are running AI-assisted reads at all.

Third: what does your PACS integration agreement say about data portability and model change notification? Most agreements say very little on either point — and most health systems discover this when a PACS transition is already underway.

03

Four Options

Option A

Continue on current clearance documentation and existing vendor contract — no independent model version monitoring

Zero operational friction. Governance gap is invisible until an adverse event or regulatory audit makes it visible. Radiologist liability and insurer non-disclosure risk remain.

Option B

Recommended

Implement model version governance: 30-day advance notice requirement, version log, local validation before production updates

Governance floor — not a full audit program, just version visibility and advance notice. Operationally achievable without renegotiating core commercial terms.

Option C

Commission a demographic performance audit: validate AI accuracy against your scanner fleet and patient population

6–12 months to complete. Closes the population validity gap but does not address model update notification or insurer disclosure. Appropriate as a follow-on, not a standalone first step.

Option D

Issue an operational review: suspend AI-assisted read workflows pending governance documentation

Operationally disruptive. Creates a clean governance baseline. Appropriate for health systems that have discovered material documentation gaps requiring a reset.

04

Recommendation

Negotiate a contract amendment requiring the vendor to: (1) provide written 30-day advance notice of any algorithm update in production, (2) supply version-specific performance data for your scanner fleet, (3) grant the health system the right to delay an update pending internal validation, and (4) maintain a model version change log accessible to your compliance team.

Notify your professional liability insurer of your current radiology AI deployment configuration before the next policy renewal. Provide the product name, FDA clearance number, deployment scope, and workflow role. This is standard risk management practice — not a concession of liability.

Commission an alert fatigue baseline study within 90 days: measure radiologist alert acknowledgment rates, time-to-review, and override rates by alert category. If no pre-deployment baseline exists, use the current period and measure again at 90 and 180 days.

Confirm in writing from the vendor the exact algorithm version running in production today. If the vendor cannot provide this in writing, escalate to CMO and general counsel — the inability itself is a material finding.

05

Risks If Not Resolved

1.

Algorithm update without notification — stale clearance documentation in production

Vendor pushes model update; hospital continues operating under original 510(k) clearance while running a materially different algorithm. Gap is invisible until a regulatory audit or adverse event investigation surfaces it.

2.

Alert fatigue degrading detection accuracy below pre-AI baseline

Alert volume grows faster than threshold calibration. Radiologists develop acknowledgment shortcuts. Detection accuracy for flagged finding classes falls below the department's pre-AI baseline — visible only in retrospective outcomes analysis.

3.

PACS vendor transition — AI data lock-in with no portability provision

PACS transition initiated. AI integration is architecturally tied to existing PACS. Vendor agreement has no API continuity or data export right. Seven years of AI-assisted read history and alert logs stays with the vendor.

4.

Malpractice claim reveals undisclosed AI deployment to insurer

Missed finding claim proceeds to discovery. Plaintiff counsel requests AI tool documentation. Health system insurer had not been notified of AI deployment configuration. Insurer initiates coverage review. Policy renewal is now framed by an undisclosed material fact.

06

Seven Questions Before Your Next Radiology AI Review

1. What algorithm version is running in your radiology AI production environment right now — and does the vendor have a written record they can provide you today?
2. Does the version currently in production match the version number in your FDA 510(k) clearance file — and when was the last time someone formally verified this?
3. What is your notification protocol when the vendor pushes a model update — and does your current contract give you the right to delay that update pending internal validation?
4. Has your professional liability insurer been notified of your radiology AI deployment configuration, and does your policy explicitly address AI-assisted clinical decision support?
5. What does your PACS integration agreement say about data portability — specifically, who owns the AI-assisted read history and alert logs at contract termination?
6. Have you measured radiologist alert acknowledgment rates and override rates since deployment — and do you have a pre-deployment baseline to compare against?
7. Which radiology AI alerts in your current configuration trigger direct workflow actions (paging, escalation) without radiologist sign-off — and has your medical staff committee formally reviewed those workflows?

AI INSIGHT LAB

The 510(k) Gap

The AI clearance your vendor holds covered the algorithm at the time of the study. The governance architecture — version monitoring, insurer disclosure, PACS data rights — has not been built at the same pace as the deployment. Building it now is cheaper than building it after an adverse event or a regulatory audit makes the gap visible.

Read the full brief at aiinsightlab.cloud/memos/radiology-ai-diagnostic-imaging