IntuitionLabs

INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

HTI-5 Source Attributes: Clinical AI Evidence Checklist

> Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-11-08 · hti-5 source attributes · clinical ai evidence checklist · predictive decision support · onc hti-5 final rule · hti-1 source attributes · clinical ai vendor evaluation · model card procurement · fda vs onc certification · clinical decision support procurement

Original article: https://intuitionlabs.ai/articles/clinical-ai-buyer-evidence-checklist-hti-5



In this article

1. Executive Summary
2. Introduction and Background
3. Which Regime Applies: A Decision Table
4. Key Changes: HTI-1 Today, HTI-5 Tomorrow
5. The Buyer's Evidence Request Matrix
6. Evaluation Worksheet: Local Validation, Bias, Updating, Drift, Monitoring
7. Implementation Considerations and Process Changes
8. Data Analysis and Evidence
9. Implications and Future Directions
10. Frequently Asked Questions (FAQs)
11. Conclusion

Executive Summary

Buyers of predictive clinical decision support now face a changed evidence baseline. The Office of the National Coordinator for Health Information Technology (ONC, within the Assistant Secretary for Technology Policy, ASTP) finalized its "HTI-5" deregulatory rule in the Federal Register on **October 28, 2026**, and the rule is effective **December 28, 2026**, apart from three amendatory instructions that start January 1, 2027 (www.federalregister.gov). For decision support interventions (DSIs), the final amendatory text removes paragraphs (b)(11)(iv) through (vi) of 45 CFR 170.315 (www.federalregister.gov). Those paragraphs carried the **source attribute** support, source attribute access and modification, and predictive-DSI **intervention risk management** requirements that HTI-1 introduced in 2024. ONC's stated reasons were that it identified no evidence of reliance on the requirements (www.federalregister.gov) and that most clinical users lack the time or technical background to use the information (www.federalregister.gov).

The practical consequence is a shift of evidence collection from a certification baseline to the buyer's contract and governance file. HTI-1 required a certified module to support up to **31** predictive-DSI source attributes (www.healthit.gov), but only for models the developer supplied (www.healthit.gov). After the effective date, ONC says the information is not "unknown or unavailable" and encourages users to ask developers for it (www.federalregister.gov). A large majority of commenters opposed the proposal (www.federalregister.gov), and, in its discussion of the legacy clinical decision support criterion, ONC said it will further consider whether the prior CDS source-attribute transparency requirements in § 170.315(a)(9)(v), or similar requirements, should be reinstated (www.federalregister.gov). That statement concerns the legacy evidence-based-intervention requirements, not the removed predictive-DSI requirements. A stable, versioned request set is therefore safer than one keyed to a single rule text; the scope bullet under "What remains uncertain" shows how far the final rule moved from the proposal.

This report supplies that request set: a decision table separating FDA device authorization from ONC certification, an eleven-row request matrix keyed to the former HTI-1 attribute categories, a missing-versus-failed scoring rubric, an evidence-gap register template, and a contract and recheck routine. The evidence base for urgency is quantitative. Hospitals reported 71% use of EHR-integrated predictive AI in 2024

(www.healthit.gov), yet among 903 FDA-listed AI devices analyzed in a 2025 peer-reviewed study, 55.9% reported a clinical performance study, and among the 505 studies reported only 8.1% were prospective (a share of studies, not of devices) (pmc.ncbi.nlm.nih.gov). FDA says its own public summaries do not include most of the submitted information (www.fda.gov).

Three cautions apply. FDA authorization and ONC certification address different scopes, and neither proves benefit at a given hospital. ONC's own test method sets no specific validity or fairness thresholds (healthit.gov). And the sources reviewed here are still mid-transition: as of November 8, 2026 the Code of Federal Regulations (eCFR) page for 170.315 was current only to November 5, 2026 (www.ecfr.gov), so readers should confirm codification status before relying on any paragraph number. This is analysis, not legal advice; compliance questions belong with counsel and regulatory affairs staff.

Introduction and Background

Hospital innovation committees, medtech product teams, and life-sciences partners routinely ask the same question when a predictive tool reaches procurement: what evidence is enough? The question became harder in 2026. A proposed rule published on December 29, 2025 would have removed the DSI transparency requirements, with comments due by 5 p.m. Eastern Time on February 27, 2026 (www.federalregister.gov). ONC's proposal described the change as removing requirements "regarding source attribute support, access, and modification" (www.federalregister.gov). The final rule followed in October 2026. Earlier coverage, including the proposal itself, describes the removal only as a proposal, and ONC's decision support test method page still displays the full HTI-1 text under Certification Companion Guide version 1.10, whose revision note reads "Removed expired standards and deadlines." (healthit.gov) This report treats the Federal Register final rule as controlling and the test method as the best available description of what the former attributes contained.

The checklist below is deliberately not an overview of the law. It asks what a supplier should be able to show, what artifact is acceptable, who inside the buyer's organization owns the review, and what to record when the evidence is absent. The row structure follows the categories ONC used for predictive DSI source attributes (purpose, development data, fairness, external validation, performance, maintenance, updating) because those categories were designed around the questions a clinician or governance committee needs answered. ONC's own stated aim for them was to let users judge whether a predictive DSI is "acceptably fair, appropriate, valid, effective, and safe." (www.federalregister.gov)

First-party perspective. IntuitionLabs is a life-sciences and AI consultancy. This report does not evaluate IntuitionLabs offerings. The material below is an analyst framework built from regulator, standards-body, and peer-reviewed sources. Where a statement is an editorial policy view, it is labeled as such.

Which Regime Applies: A Decision Table

Three different authorities touch a predictive clinical tool, and they do not overlap neatly. FDA regulates devices. ONC regulates health IT developers that voluntarily certify modules under its Certification Program. A third source is the Department of Health and Human Services (HHS) nondiscrimination rule, which places a duty on covered entities that use patient care decision support tools. ONC itself noted in HTI-1 that FDA and

ONC "have separate and distinct authorities and regulate for separate and distinct purposes." (www.federalregister.gov) ONC also declined to exclude FDA-regulated or non-device software from its predictive DSI definition (www.federalregister.gov), so a single tool can sit inside both regimes.

On the FDA side, the Clinical Decision Support (CDS) Software guidance, reissued January 29, 2026 (www.fda.gov), keeps the four-criterion test for non-device software. Criterion 4 requires that the clinician be able to independently review the basis for a recommendation (www.fda.gov), and FDA does not consider functions intended for a critical, time-sensitive task or decision to meet it (www.fda.gov). Its example of software that detects a life-threatening condition such as stroke or sepsis and alerts a clinician is a device function (www.fda.gov). Software aimed at patients or caregivers is addressed separately and treated as meeting the device definition (www.fda.gov).

On the ONC side, a predictive DSI is technology based on algorithms or models that derive relationships from training data (www.ecfr.gov). The Certification Companion Guide says a developer is not accountable for populating source attributes for predictive DSIs it does not supply (healthit.gov), while "supplied" includes other-party models the developer has included, because it then has taken on stewardship and accountability (healthit.gov). *Table 1* below organizes these overlaps into a decision table. It is an analytical aid, not a classification of any product.

Table 1. Applicability decision table (analyst framework, status as of November 8, 2026)

Scenario	FDA device authorization	ONC certified health IT obligation	What the buyer should do
A. Predictive alert inside an EHR module, supplied by the EHR developer, time-critical use	Likely a device function under the CDS guidance. Ask for the authorization number and indications for use.	HTI-1 source attributes applied to supplied predictive DSIs (healthit.gov) until HTI-5 takes effect on December 28, 2026. After that, no federal attribute baseline.	Request both the FDA documents and the former attribute set by contract.
B. Risk score inside an EHR module, supplied, clinician-reviewable, not time-critical	May be non-device CDS if all four criteria hold. Ask the supplier for its written criteria analysis.	Same as A.	Request the attribute set. Do not treat "non-device" as a signal of lower evidence needs.
C. Third-party model connected to a certified EHR but not supplied by the developer	Depends on the model's own classification.	Developer is not accountable for populating attributes for models it does not supply. Developers must give users the capability to populate attributes (HTI-1 text).	Obtain the evidence from the model owner directly. Record who supplies the model.
D. In-house model built by the hospital	Classification depends on the facts of development and use. Confirm with regulatory counsel.	Outside the developer's duty.	The hospital is its own supplier: apply the same rows to internal teams.
E. Standalone clinical AI product, not part of a certified module	Often a device with a public authorization record (www.fda.gov).	No ONC certification duty unless the product is part of a certified module.	Use FDA documents plus Table 2 rows. Ask what the FDA summary omits.

Scenario	FDA device authorization	ONC certified health IT obligation	What the buyer should do
F. Patient-facing or caregiver-facing recommendation tool	Treated as meeting the device definition under the CDS guidance.	Not specific to ONC unless integrated.	Request authorization status and usability evidence for lay users.

Table 1 shows why a single "FDA cleared" or "ONC certified" label cannot settle diligence. Rows A and B can look identical to a clinician at the bedside yet carry different FDA status, and rows C and D fall outside the developer's certification duty entirely. In a Regulatory Review opinion piece published August 6, 2026, a clinical AI founder argued that non-device decision support, internally developed models, and generative AI assistants frequently fall outside FDA oversight (www.theregreview.org). That is an editorial position, not a measurement, but the classification logic in the table is consistent with it. For federal nondiscrimination purposes, the HHS definition of "patient care decision support tool" covers any automated or non-automated tool, mechanism, method, technology, or combination thereof (www.federalregister.gov), so the duty reaches beyond AI.

Key Changes: HTI-1 Today, HTI-5 Tomorrow

What the HTI-1 criterion required

HTI-1, published January 9, 2024 and effective February 8, 2024, created the DSI criterion at 170.315(b)(11). ONC's fact sheet states that it supports 31 source attributes for predictive DSIs. Developers with the earlier clinical decision support certification had to update to the new criterion by December 31, 2024. Conformance is by attestation to the certification body (healthit.gov), not by an ONC-run performance test.

The predictive attribute set grouped into categories that map directly to a buyer's questions:

- **Details and output:** developer name and contact information.
- **Purpose:** intended use (healthit.gov) and intended decision-making role.
- **Cautioned out-of-scope use:** tasks, situations, or populations where users are cautioned against use (healthit.gov), plus known risks and limitations.
- **Development data:** inclusion and exclusion criteria for training data (healthit.gov), demographic representativeness (healthit.gov), and relevance of training data to the deployed setting.
- **Fairness in development:** approaches to manage, reduce, or eliminate bias.
- **External validation:** the testing party and the validation process (healthit.gov).
- **Performance:** validity on same-source test data, fairness on external data, and outcome-evaluation references.
- **Ongoing maintenance and updating:** validity in local data (healthit.gov), fairness in local data, update process and frequency (healthit.gov).

Separately, HTI-1 required risk management practices for supplied predictive DSIs, covering risk analysis across validity, reliability, robustness, fairness, intelligibility, safety, security, and privacy (healthit.gov), mitigation, and governance of data acquisition and use. ONC did not prescribe a framework. For several attributes the module had to indicate when information was not available for review (healthit.gov), and ONC

acknowledged that models drawn from the literature may have training data that is not available nor re-creatable. This last point matters to buyers: "unknown" under HTI-1 was a permitted answer, and was never evidence of absence.

What HTI-5 removes

The final rule removes paragraphs (b)(11)(iv) through (vi), which ONC ties to "model cards" that it says "essentially function as nutrition labels" (www.federalregister.gov). It states, "We have finalized removal of the requirements related to" model cards (www.federalregister.gov). The final text also removes the transition-of-care intervention-enablement paragraph (b)(11)(ii)(B) and the conforming maintenance-of-certification duty to review and update source attributes and risk management practices (www.federalregister.gov). The legacy clinical decision support criterion (a)(9) is removed and reserved, with ONC noting this has no further substantive effect because the criterion had already lapsed. ONC's chart lists the DSI criterion as "Revise," not "Remove," with the effective date of the final rule (healthit.gov).

What HTI-5 retains and what ONC says about voluntary information

The electronic feedback requirement for evidence-based DSIs stays, redesignated (www.federalregister.gov). ONC states that removal "does not mean that such information about intended use, provenance, training data, limitations, performance, and known risks is unknown or unavailable" (www.federalregister.gov), and it encourages users to request it from developers. It adds that risk management for supplied predictive DSIs "may still be considered best practices", and that industry-led transparency efforts "have promise". The rule does not require developers to keep publishing the information. Law-firm commentary characterizes the choice as voluntary: one firm advises developers and providers to determine whether to continue providing model card information and risk management voluntarily (www.mcdermottlaw.com), while another observes that vendors continuing to publish model cards may find it sets them apart in procurement (www.reedsmith.com). Hospital and clinician commenters objected, according to a third firm, that the change shifts the burden of vetting AI onto providers (www.hklaw.com).

What remains uncertain

- **Codification lag.** The eCFR page for 170.315 was current to November 5, 2026 (www.ecfr.gov), before the December 28 effective date.
- **Stale agency pages.** The test method page was last revised under version 1.10 and, per this review's reading of the page text, does not mention HTI-5.
- **Future rulemaking.** In its discussion of the legacy clinical decision support criterion, ONC said it will further consider whether the prior CDS source-attribute transparency requirements in § 170.315(a)(9)(v), or similar requirements, should be reinstated. This statement does not address the removed predictive-DSI requirements in (b)(11)(iv)-(vi).
- **Commenter sentiment versus outcome.** Many commenters said removal was premature because the requirements had been in effect for less than two years, yet the provision was finalized.
- **Scope of criteria changes.** The proposal would have removed 34 of 60 criteria (www.federalregister.gov), and the final rule removed 14 and revised seven (www.federalregister.gov), retaining 39. Scope moved materially between proposal and final text, which is itself a reason to build

checklists that do not depend on one rule version.

The Buyer's Evidence Request Matrix

The matrix below is the core deliverable. Each row pairs a former ONC attribute category with a request, an acceptable artifact, a suggested internal owner, and red flags. Rows 1 to 8 correspond to attribute categories; rows 9 to 11 add governance and lifecycle items drawn from the former risk management practices and from FDA, Joint Commission, and health-sector guidance.

Two distinctions keep the matrix honest. First, a certified-module obligation (what a certified developer had to support for supplied models) differs from a voluntary request to a third-party model owner, so the "basis" column below labels each row accordingly. Second, FDA public summaries are a partial view: FDA states the list is not a comprehensive resource of AI-enabled devices (www.fda.gov), with inclusion based primarily on AI-related terms in the summary descriptions (www.fda.gov).

Table 2. Evidence request matrix keyed to former HTI-1 source attribute categories

#	Evidence domain	ONC anchor (former HTI-1 attribute)	Request to supplier	Accepted artifact	Suggested owner	Red flags
1	**Intended use and role**	Intended use and decision-making role	State the clinical question, user, setting, and whether output informs, augments, or replaces judgment.	Intended-use statement matching FDA indications (if authorized) and the labeling.	Clinical lead	Marketing claims broader than the written intended use.
2	**Intended population and out-of-scope use**	Cautioned out-of-scope use and known limitations	Name excluded populations, settings, and tasks.	Labeling, model card, or instructions for use with explicit exclusions.	Clinical lead and patient safety	No exclusions listed for any tool.
3	**Data provenance**	Inclusion and exclusion criteria for training data	Describe sources, time window, sites, and label definitions.	Data sheet with collection dates, site counts, and exclusion logic.	Data science lead	"Proprietary" with no summary; labels derived from billing codes with no validation.
4	**Representativeness and relevance**	Demographic representativeness and relevance to deployed setting	Provide demographic breakdown of development data versus your population.	Table of age, sex, race and ethnicity, and payer mix against the local population.	Health equity lead	Only aggregate counts; demographic fields missing.

#	Evidence domain	ONC anchor (former HTI-1 attribute)	Request to supplier	Accepted artifact	Suggested owner	Red flags
5	**External validation**	Testing party and validation process	Name who tested on which independent sites and when.	Peer-reviewed paper or independent report. FDA guidance suggests test data from sites different from development sites (www.fda.gov).	Chief medical informatics officer	Validation only on the development site; supplier-authored only.
6	**Performance and fairness metrics**	Same-source validity and external fairness	Provide discrimination, calibration, and subgroup metrics with confidence intervals.	Complete results tables, ideally following a reporting checklist such as TRIPOD+AI (pubmed.ncbi.nlm.nih.gov).	Biostatistician	Single headline metric; no subgroup results; "no notable bias" with no data.
7	**Outcome and workflow evidence**	Outcome-evaluation references	Provide evidence of effect on decisions and patient outcomes in live use.	Prospective or live-use study; early-stage reporting aligned with DECIDE-AI (www.bmj.com).	Clinical lead	Retrospective-only evidence presented as outcome proof.
8	**Local validation plan**	Validity and fairness in local data	Will the supplier support silent-mode local validation before go-live?	Written protocol, data-access plan, pre-set acceptance thresholds.	Data science lead	Supplier declines silent validation; no sandbox.
9	**Monitoring and drift**	Maintenance and monitoring (former maintenance category)	Specify metrics tracked, frequency, thresholds, and who is notified.	Monitoring plan; sample dashboard; alerting service levels.	Quality and safety	Monitoring "by customer" with no tools provided.
10	**Updates and change control**	Update process and frequency and corrective frequency	Describe retraining triggers, notice periods, version labeling, rollback.	Change-control policy; for FDA devices, the authorized change plan (www.fda.gov).	IT and procurement	Silent model changes; no version identifier visible to users.
11	**Risk management and incident handling**	Risk analysis, mitigation, governance	Provide risk register, escalation path, and event-reporting terms.	Documented risk process; event-notification clause; contact roster.	Patient safety and legal	No named contact; no defined timeline for notice.

Table 2 should be read left to right as a negotiation script. Rows 1 to 4 define whether the tool was built for the buyer's patients, rows 5 to 8 test whether the claim survived contact with different data, and rows 9 to 11 govern what happens after go-live. Where HTI-1 imposed a certification-level duty, the "ONC anchor" now serves as vocabulary only; where the supplier is a third-party model owner, the request was always voluntary. The matrix aligns with sector guidance on vendor questions: the Joint Commission and Coalition for Health AI (CHAI) guidance advises asking whether vendors are willing to tune and validate a sample representative of the deployment setting (digitalassets.jointcommission.org), and a vendor-disclosure framework from the Health AI Partnership exists because hospitals have had limited guidance on what information to request from vendors during procurement (ai.nejm.org).

Evaluation Worksheet: Local Validation, Bias, Updating, Drift, Monitoring

Scoring rubric: missing versus failed

A common error in diligence files is treating absent evidence and adverse evidence as the same entry. They carry different actions. The rubric below uses two independent fields per row: an **evidence level** and a **result status**. No numerical vendor scores are supplied here, because scoring requires the buyer's own thresholds.

- **Evidence level 0, Missing:** nothing was provided, or the supplier answered "unknown." Action: request again, then record as a gap with a compensating control.
- **Evidence level 1, Asserted:** a statement without an artifact. Action: request the artifact.
- **Evidence level 2, Documented:** an artifact with method, population, metrics, and date, supplier-authored.
- **Evidence level 3, Verified:** independently authored (peer-reviewed or third-party) or reproduced on the buyer's local data.
- **Result status "Met":** evidence meets thresholds the committee set before opening the supplier's data.
- **Result status "Failed":** evidence was provided and falls below those pre-set thresholds, or shows a subgroup disparity beyond tolerance. Action: do not deploy for the affected scope without a documented exception.
- **Result status "Not assessable":** no evidence was provided (Level 0), or evidence exists but cannot be compared to the local population.

Under this scheme, a row can be "Level 0, not assessable" (a gap) or "Level 3, Failed" (a rejection). Both are important, but only the second is a finding about the tool. HTI-1 itself allowed "unknown" for several attributes, which is why a blank is not a failure.

Local validation

Evidence shows why local checks matter. A review of 516 diagnostic imaging studies found only 6% performed external validation (pmc.ncbi.nlm.nih.gov). A systematic review of deep-learning radiology algorithms found 81% (70 of 86) reported at least some decrease in performance externally compared with internal results (pubmed.ncbi.nlm.nih.gov). A New England Journal of Medicine perspective defines dataset shift as underperformance arising from a mismatch between the development data and the deployment data (pmc.ncbi.nlm.nih.gov) and recommends running models silently first (pmc.ncbi.nlm.nih.gov).

- Set accuracy, calibration, and subgroup tolerances before seeing results.
- Run the tool silently on a recent local cohort; compare against the supplier's reported metrics.
- Record the sample window, size, and exclusions in the register.
- Treat a vendor's refusal to enable silent mode as a Level 0 finding on row 8.

Bias and fairness

The federal duty on covered entities is separate from the supplier's attributes. The nondiscrimination rule at 45 CFR 92.210 imposes an ongoing duty to make reasonable efforts to identify uses of patient care decision support tools that use protected characteristics as input variables (www.ecfr.gov) and to mitigate the risk of discrimination resulting from those tools (www.ecfr.gov). An HHS Office for Civil Rights letter dates those obligations to May 1, 2025 (www.hhs.gov). Practical steps: ask which input variables encode race, sex, age, or disability; ask whether subgroup performance was reported; and compare to the local population. Public FDA-authorization summaries are weak on this point: one study of 692 summaries found only 3.6% reported race and ethnicity and 99.1% provided no socioeconomic data (www.nature.com), and 81.6% did not report the age of study subjects (www.nature.com). FDA's draft lifecycle guidance warns that acceptable performance in certain subgroups may mask lower performance in others (www.fda.gov).

Updating and change control

FDA's final guidance on predetermined change control plans, originally issued December 4, 2024 (www.fda.gov), expects a modification protocol with data management, re-training, performance evaluation, and update procedures, with test data separate and independent from development (www.fda.gov). It also asks whether modifications differ across clinical sites (www.fda.gov) and recommends labeling disclose an authorized plan (www.fda.gov). A buyer can use these headings even for non-device tools. The AMA's AI evaluation guide flags, as an area to watch, any absence of description of how model updates are managed (www.ama-assn.org).

Drift and monitoring

FDA's draft lifecycle guidance, issued January 7, 2025 (www.fda.gov), says performance may change or degrade over time (www.fda.gov) and lists monitoring inputs such as changes in patient demographics or disease prevalence (www.fda.gov). Clinical-model research shows the pattern: in a Veterans Affairs kidney-injury study, calibration declined as models increasingly overpredicted risk (pubmed.ncbi.nlm.nih.gov), and the authors argued updating should respond to rapid drift instead of fixed annual or biannual schedules (pubmed.ncbi.nlm.nih.gov). The National Institute of Standards and Technology (NIST) AI Risk Management Framework notes AI systems may require more frequent maintenance and triggers for conducting it (nvlpubs.nist.gov). The Joint Commission and CHAI guidance asks for feedback channels between third-party vendors and local monitors (digitalassets.jointcommission.org) and recommends that tools which inform or drive clinical decisions should be checked more often than administrative ones (digitalassets.jointcommission.org).

- Define metrics, thresholds, review cadence, and the named reviewer.
- Require version identifiers in every logged prediction.
- Pre-agree what triggers a pause, rollback, or revalidation.
- Review monitoring output at the governance committee, not only inside the supplier dashboard.

The evidence-gap register

Table 3 is a template. It contains no real suppliers and no scores. The rows are **(Hypothetical Example)** placeholders showing how entries are written.

Table 3. Evidence-gap register template (Hypothetical Example rows; version 1.0)

Gap ID	Tool and version	Matrix row	Evidence level (0 to 3)	Result status	Risk if unresolved	Compensating control	Owner	Recheck trigger
G-001 (Hypothetical Example)	[Tool A, v1.0]	4 Representativeness	0	Not assessable	Performance in the local population unknown	Silent-mode local validation before go-live	Health equity lead	Supplier delivers demographic table, or 60 days elapse
G-002 (Hypothetical Example)	[Tool B, v2.3]	10 Updates	1	Not assessable	Unannounced model changes alter alert volume	Contract notice clause; version logging	IT and procurement	Any new version string
G-003 (Hypothetical Example)	[Tool C, v1.1]	6 Metrics	3	Failed	Subgroup calibration below pre-set tolerance	Restrict scope or decline	Clinical lead	New validation data or recalibration
G-004 (Hypothetical Example)	[Tool D, v1.0]	11 Incident handling	2	Met	None identified	Annual review	Patient safety	Contract renewal

Table 3 is meant to be retained as a living risk register entry per tool, not a one-time form. The recheck trigger column links register rows to the regulatory recheck routine in the next section, and the version field lets auditors see which edition of the request set was used for a given decision.

Implementation Considerations and Process Changes

Contract and change-control questions

Joint Commission and CHAI guidance says monitoring responsibility between the hospital and vendor should be settled as part of third-party procurement and contracting (digitalassets.jointcommission.org), and that data use agreements should reserve audit rights (digitalassets.jointcommission.org). A health-sector coordinating council guide includes 69 vendor risk questions in an appendix (healthsectorcouncil.org) and sample language requiring vendors to implement continuous or periodic monitoring for model drift (healthsectorcouncil.org). Combining these sources, a procurement committee can ask:

- **Disclosure continuity:** will the supplier keep providing the former source attribute information after December 28, 2026, and for how long? ONC encourages users to ask.
- **Model card cadence:** will updated cards be reissued on material changes? CHAI's Applied Model Card documentation recommends that implementers ask for updated cards annually (mc.chai.org).

- **Notice of change:** how many days of advance notice precede a model update, and what version identifier appears in logs?
- **Local access:** will the supplier grant data access and a test environment for silent validation?
- **Event reporting:** how does the supplier receive and respond to reports of unsafe outputs? The Joint Commission and CHAI guidance recommends treating AI near-misses and unsafe recommendations like patient safety events (digitalassets.jointcommission.org).
- **Audit and exit:** what audit rights, data return, and rollback commitments exist?
- **Regulatory status warranties:** will the supplier notify the buyer of changes in FDA status or ONC certification status?

Governance structures

ASTP data show that hospital governance is distributed: 74% of hospitals using predictive AI said multiple entities were accountable for evaluating it, and the most cited evaluator was a dedicated committee or task force for predictive AI at 66% (www.healthit.gov). The Joint Commission and CHAI guidance lists seven elements, including voluntary, blinded reporting of AI safety-related events (digitalassets.jointcommission.org), and states that it does not validate the effectiveness of AI tools themselves (digitalassets.jointcommission.org). The Joint Commission's 2026 certification program likewise does not validate or certify individual AI products (www.jointcommission.org). In other words, organizational certification is not product evidence.

Periodic regulatory-status recheck

Because the rule status moved from proposal to final within ten months, the register needs a recheck routine. A suggested cadence:

- **Now through December 28, 2026:** confirm the effective date and whether the HTI-1 text remains in force for current certified modules.
- **At each quarter:** check the Federal Register, eCFR currency date, ONC's HTI-5 final rule page, and the test method page for updates.
- **At each contract renewal:** reconfirm FDA status, including the date of the current CDS guidance.
- **On any new ONC proposal:** reassess the matrix, because the request set should not be keyed to a single rule text and the scope of ONC's criteria changed between the proposed and final rules (see "Scope of criteria changes").

Data Analysis and Evidence

The numbers below are the quantitative basis for treating public regulatory status as insufficient. They come from regulators and peer-reviewed analyses, and each carries the denominator and date the source states.

Adoption and evaluation in US hospitals (ASTP, 2024 AHA IT Supplement). In 2024, 71% of non-federal acute care hospitals used predictive AI integrated with the EHR. Of hospitals using it, 82% evaluated for accuracy, 74% for bias, and 79% conducted post-implementation evaluation or monitoring (www.healthit.gov), but 15% answered "don't know" on accuracy evaluation, 21% on bias, and 18% on monitoring, and those answers were excluded from the denominators, so the shares may overstate practice (www.healthit.gov). Use was 86% at system-affiliated hospitals versus 37% at independent hospitals, and 80% used predictive AI from

their EHR developer while 52% used third-party AI (www.healthit.gov). The survey had a 51% response rate (N = 2,253). ASTP itself flags demonstrating real-world performance as a persistent implementation challenge. An earlier analysis of 2023 data found that hospitals with their own models, high operating margins, or system membership were more likely to evaluate locally, which the authors describe as a risk of a new digital divide (www.healthaffairs.org).

FDA-authorized AI devices. FDA states it has authorized over 1,600 AI-enabled devices as of September 2026. A tally of the downloadable FDA list by this review, covering 1,614 rows with a newest decision date of June 29, 2026, gives 226 final decisions in 2023, 235 in 2024, 335 in 2025, and 181 in 2026 through June 29. This is an approximate analyst count, not independently reproduced, of a list FDA says it updates periodically, so figures will differ from other published counts depending on list vintage. Source file: www.fda.gov.

Public evidence in device summaries. Because buyers often rely on these summaries, the reporting gaps are decision-relevant:

- **Clinical performance studies reported:** 505 of 903 devices (55.9%) (pmc.ncbi.nlm.nih.gov); 218 (24.1%) explicitly stated no performance study (pmc.ncbi.nlm.nih.gov).
- **Study type among the 505 reported studies (share of studies, not of devices):** 8.1% prospective and 2.4% randomized (pmc.ncbi.nlm.nih.gov).
- **Randomized or prospective data in a separate review of 691 clearances:** 1.6% randomized and 7.7% prospective (jamanetwork.com).
- **Registered trials:** 34 of 1,357 devices (2.5%) had registered clinical trials, and 3 (0.2%) reported patient outcomes (journals.plos.org).
- **Post-market:** only 9.0% contained a prospective study for post-market surveillance (www.nature.com).

These studies measure what public documents say, not what manufacturers know; the authors of the first study note that detailed descriptions of development were absent from most publicly provided summaries (pmc.ncbi.nlm.nih.gov). FDA's own public comment request acknowledges that many AI devices are evaluated primarily through retrospective testing or static benchmarks (www.fda.gov), which are not designed to predict behavior in dynamic, real-world environments (www.fda.gov).

Regulatory cost estimates. ONC counted 251 certified products from 206 developers for the DSI criterion and estimated annual savings from the DSI change of $7,361,579 to $53,275,503 (www.federalregister.gov). These are ONC's developer-side burden estimates. These figures describe developer-side savings only; this review found no published estimate of the diligence cost that shifts to buyers.

Interpretation. Availability of a clearance record does not equal availability of evidence. A 510(k) is a premarket submission made to demonstrate that a device is as safe and effective as a legally marketed device (www.fda.gov); FDA describes its listed devices as having met premarket review of overall safety and effectiveness, including evaluation of study appropriateness (www.fda.gov). Neither statement says the device performs as expected at a specific hospital. For a published reporting standard, TRIPOD+AI offers a 27-item checklist (pubmed.ncbi.nlm.nih.gov), and a risk-of-bias tool, PROBAST+AI, uses 18 targeted signalling questions for evaluating models (pubmed.ncbi.nlm.nih.gov).

Implications and Future Directions

For hospital committees. The removal of a federal attribute baseline changes where the evidence lives. An editorial view published before the final rule argued that removing the floor does not relieve a burden but shifts it onto the buyer (www.theregreview.org). The final rule's text is consistent with that direction of travel: ONC pointed users to their developers for the information. Committees with lean staff, particularly independent and small hospitals where predictive AI use is lower, may find that a standard request set reduces the cost of each evaluation.

For medtech and product teams. Suppliers that already maintain structured transparency documents have a reusable asset. CHAI's Applied Model Card is described as a standardized template usable in procurement processes (mc.chai.org). FDA-authorized products can reuse change-plan and monitoring documentation. Suppliers without such documents face repeated bespoke questionnaires.

Open questions.

- Whether ONC acts on its stated plan to consider reinstating the legacy CDS source-attribute transparency requirements in § 170.315(a)(9)(v), and whether it later revisits predictive-DSI transparency.
- Whether state or accreditation programs fill the gap. The NIST framework is voluntary (www.nist.gov) and Joint Commission certification is organizational (www.jointcommission.org).
- Whether suppliers continue voluntary disclosure; this report found no measured data yet on that behavior.
- Whether evidence-gap registers improve outcomes. No measured outcome data exist; the register is a risk-management practice, not a proven intervention.

Frequently Asked Questions (FAQs)

What is HTI-5 and how does it relate to ONC certification? HTI-5 is the ONC rule "Health Data, Technology, and Interoperability: ONC Deregulatory Actions to Unleash Prosperity," published October 28, 2026. It resets the scope of the Certification Program and reduces developer compliance costs. The criteria apply to certified health IT modules supplied by developers, not directly to hospitals.

What did the HTI-5 proposed rule say about decision support interventions? It proposed removing source attribute support, access, and modification requirements and said ONC had no publicly available evidence that a single doctor, nurse, or administrator had accessed, recorded, or modified a source attribute (www.federalregister.gov).

What are predictive DSI source attributes? They are the descriptive fields (purpose, training data, fairness approaches, validation, performance, maintenance, updating) that a certified module had to support for supplied predictive DSIs under HTI-1. See the matrix in Table 2.

Do HTI-1 source attribute rules still apply? As of November 8, 2026, the HTI-5 changes take effect December 28, 2026, and the eCFR page reviewed for this report was current only to November 5 (www.ecfr.gov). Confirm current status with counsel.

How do AI model card requirements fit health IT certification? ONC described the source attribute requirements as supporting model cards, and removed them. Voluntary model cards continue as an industry practice (mc.chai.org).

How should a health system evaluate clinical AI vendors? Use the matrix in Table 2, the rubric in the worksheet, and the register in Table 3, then confirm contract terms for monitoring, notice, and audit (digitalassets.jointcommission.org).

What questions belong in clinical decision support AI procurement? See the contract and change-control list; the AMA guide adds recent-data validation (www.ama-assn.org) and cautions about tools reporting only retrospective performance (www.ama-assn.org).

Does FDA clearance prove local clinical benefit? No. Clearance is a premarket determination for an intended use. Local validation remains necessary.

Conclusion

The central finding is that evidence sufficiency, not regulatory labeling, should drive clinical AI purchasing. ONC's HTI-5 final rule, published October 28, 2026 and effective December 28, 2026, removes the certified-module requirement to support source attributes and predictive-DSI risk management. FDA authorization addresses a different scope and publicly documents only part of what was submitted. Neither regime sets validity or fairness thresholds for a given hospital's patients.

A versioned request set based on the former ONC attribute categories lets a buyer ask for the same evidence whether HTI-1 text, HTI-5 text, or some future ONC action governs. The elements are a decision table that separates scopes, an eleven-row matrix, a rubric that distinguishes missing from failed evidence, a register with recheck triggers, and a contract list covering notice, monitoring, and audit. Each should be treated as a working document, revised when the regulatory status, the supplier's model version, or the local population changes.

Readers should confirm the live status of the rule, the codified paragraph text, and any product-specific FDA classification with qualified counsel and regulatory staff before acting on this framework.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://www.federalregister.gov/d/2026-21943
2. https://www.healthit.gov/sites/default/files/page/2023-12/HTI-1_DSI_fact%20sheet_508.pdf
3. https://www.healthit.gov/sites/default/files/2025-09/hospital-trends-use-evaluation-and-governance-predictive-ai-2023-2024.pdf
4. https://pmc.ncbi.nlm.nih.gov/articles/PMC12044510
5. https://www.fda.gov/medical-devices/artificial-intelligence-enabled-medical-devices/list-artificial-intelligence-enabled-medical-devices
6. https://healthit.gov/test-method/decision-support-interventions
7. https://www.ecfr.gov/current/title-45/subtitle-A/subchapter-D/part-170/subpart-C/section-170.315

8. https://www.federalregister.gov/documents/2025/12/29/2025-23896/health-data-technology-and-interoperability-astponc-deregulatory-actions-to-unleash-prosperity
9. https://www.federalregister.gov/documents/2024/01/09/2023-28857/health-data-technology-and-interoperability-certification-program-updates-algorithm-transparency-and
10. https://www.fda.gov/media/109618/download
11. https://www.ecfr.gov/current/title-45/section-170.102
12. https://www.fda.gov/medical-devices/digital-health-center-excellence/artificial-intelligence-enabled-medical-devices
13. https://www.theregreview.org/2026/08/06/craige-the-governance-gap-in-clinical-ai
14. https://www.federalregister.gov/documents/2024/05/06/2024-08711/nondiscrimination-in-health-programs-and-activities
15. https://healthit.gov/resources/hti-5-final-rule-removals-or-revisions-of-certification-criteria
16. https://www.mcdermottlaw.com?p=50392216
17. https://www.reedsmith.com/our-insights/blogs/health-industry-washington-watch/102o6vt/onc-finalizes-hti-5-de-regulating-less-than-anticipated-and-punting-on-informati
18. https://www.hklaw.com/en/insights/publications/2026/10/hti5-promotes-deregulatory-certification-actions
19. https://www.fda.gov/media/184856/download
20. https://pubmed.ncbi.nlm.nih.gov/38626948
21. https://www.bmj.com/content/377/bmj-2022-070904
22. https://www.fda.gov/media/166704/download
23. https://digitalassets.jointcommission.org/api/public/content/dcfcf4f1a0cc45cdb526b3cb034c68c2
24. https://ai.nejm.org/doi/full/10.1056/AIp2500985
25. https://pmc.ncbi.nlm.nih.gov/articles/PMC6389801
26. https://pubmed.ncbi.nlm.nih.gov/35652114
27. https://pmc.ncbi.nlm.nih.gov/articles/PMC8665481
28. https://www.ecfr.gov/api/renderer/v1/content/enhanced/current/title-45?part=92&section=92.210
29. https://www.hhs.gov/sites/default/files/ocr-dcl-section-504-section-1557-disability.pdf
30. https://www.nature.com/articles/s41746-024-01270-x
31. https://www.ama-assn.org/system/files/ai-tool-evaluation-guide.pdf
32. https://pubmed.ncbi.nlm.nih.gov/28379439
33. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf
34. https://healthsectorcouncil.org/wp-content/uploads/2026/04/AI-Third-Party-Risk-Guide.pdf
35. https://mc.chai.org/v0.1/documentation.pdf
36. https://www.jointcommission.org/en-us/knowledge-library/news/2026-05-responsible-use-of-ai-in-healthcare-certification
37. https://www.healthaffairs.org/doi/10.1377/hlthaff.2024.00842
38. https://www.fda.gov/media/178541/download?attachment
39. https://jamanetwork.com/journals/jama-health-forum/fullarticle/2839236
40. https://journals.plos.org/digitalhealth/article?id=10.1371/journal.pdig.0001597
41. https://www.fda.gov/medical-devices/digital-health-center-excellence/request-public-comment-measuring-and-evaluating-artificial-intelligence-enabled-medical-device

42. https://www.fda.gov/medical-devices/premarket-submissions-selecting-and-preparing-correct-submission/premarket-notification-510k
43. https://pubmed.ncbi.nlm.nih.gov/40127903
44. https://www.nist.gov/itl/ai-risk-management-framework

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore AI enablement, pharma and biotech software development, data engineering and BI, and Veeva services. Contact IntuitionLabs to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: https://intuitionlabs.ai

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.