

REGULATION

Regulating AI in certified EHR technology: what actually binds a vendor

ABSTRACT

An EMR that ships a predictive feature meets three federal regimes and a growing state layer: health IT certification, FDA device regulation, and a voluntary assurance layer that is not regulation at all. This paper sets out what each requires, with exact citations and dates. The certification criterion at 45 C.F.R. § 170.315(b)(11) compels disclosure and documented risk management. It sets no performance threshold, and neither does anything else that reaches an EMR.

Author Naomi Hart, M.D. **Type** Regulatory review **References** 28 **Reading time** 15 min

Last reviewed September 2026

1 Three regimes, one product

An EMR vendor ships a feature that predicts something: a readmission score, a no-show probability, a suggested diagnosis code, a note drafted for signature. What law applies?

Three federal regimes reach that feature, for different reasons and with different force. The first is the ONC Health IT Certification Program, administered by ASTP/ONC, which since the HTI-1 final rule has conditioned certification on a transparency criterion for decision support,¹ codified at 45 C.F.R. § 170.315(b)(11).² The second is FDA device regulation, which asks whether the software function is a device at all, a question governed by a statutory exclusion added by the 21st Century Cures Act.³ The third is not regulation: voluntary frameworks and reporting guidelines from journals, professional bodies and an industry coalition. A fourth layer, state statute, has grown quickly since 2024 and lands mostly on providers and payers.

These are commonly conflated, and the conflation runs one way: readers assume more is required of a model's performance than actually is. What follows sets out what each regime requires, and what it does not.

31

source attributes a Predictive DSI must surface, in nine categories

34 of 60

certification criteria proposed for removal in HTI-5

0.4%

of FDA-authorized AI devices take EHR-based inputs

2 The Decision Support Interventions criterion

The operative instrument is HTI-1, *Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing*, 89 Fed. Reg. 1192–1438 (Jan. 9, 2024), RIN 0955-AA03, effective February 8, 2024.¹ It retired the legacy clinical decision support criterion at § 170.315(a)(9) and replaced it with the Decision Support Interventions criterion at § 170.315(b)(11), the first federal transparency requirement aimed at predictive algorithms in certified health IT.

The compliance date has two halves. Developers had to be certified to (b)(11) by **December 31, 2024**. Effective **January 1, 2025**, (a)(9) was removed and (b)(11) became the operative criterion for the definition of certified electronic health record technology used in CMS programs.¹

Two classes of intervention

An evidence-based DSI must expose **13 source attributes** under § 170.315(b)(11)(iv)(A). A Predictive DSI must expose **31 source attributes**, organized under **nine categories**, at § 170.315(b)(11)(iv)(B).^{2,4} ASTP/ONC defines a Predictive DSI as “technology that supports decision-making based on algorithms or models that derive relationships from training data and then produce an output that results in prediction, classification, recommendation, evaluation, or analysis.”⁴

TABLE 1

The nine source-attribute categories for Predictive DSIs, verbatim from 45 C.F.R. § 170.315(b)(11)(iv)(B). Thirty-one individual attributes sit beneath these nine headings; the headings are not themselves the attributes.

#	Category	What it discloses
1	Details and output of the intervention	What the intervention is and what it emits
2	Purpose of the intervention	The intended use
3	Cautioned out-of-scope use of the intervention	Uses the developer warns against
4	Intervention development details and input features	How it was built and what it consumes
5	Process used to ensure fairness in development	The process, not its result
6	External validation process	Whether and how it was validated elsewhere
7	Quantitative measures of performance	The numbers, with no floor set on them
8	Ongoing maintenance of intervention implementation and use	How it is kept working in the field
9	Update and continued validation or fairness assessment schedule	When it is revisited

Access, and the visible blank

Paragraph (v) requires that a limited set of identified users be able to access, record and modify source attributes, and be able to see when a source attribute value is unavailable.² The rule contemplates missing attributes and requires the gap to be shown.

Intervention Risk Management

Paragraph (vi) imposes Intervention Risk Management on Predictive DSIs the developer *supplies*. IRM has three practice areas: **risk analysis**, **risk mitigation** for identified risks, and **governance** covering how data are acquired, managed and used. The risk analysis runs across eight named dimensions: validity, reliability, robustness, fairness, intelligibility, safety, security and privacy. A plain-language summary of the practices must also be publicly available, the hyperlink submitted to the developer's ONC-Authorized Certification Body and the summary reachable without preconditions: no fee, no account, no click-through agreement.²

The supply limitation matters. Source attributes and IRM attach only to Predictive DSIs the developer supplies; for third-party or customer-configured interventions the module need only *enable* identified users to record, change and access the attributes.² Accounts claiming that HTI-1 forces a vendor to publish model cards for every third-party tool in its marketplace are wrong.

A third date is often missed. HTI-2 added privacy and security certification requirements for modules certified to § 170.315(b)(11), with a compliance date of **January 1, 2028**.⁵

3 What the criterion does not do

Nothing in § 170.315(b)(11) requires a model to be accurate. ASTP/ONC does not evaluate model performance, sets no accuracy or fairness threshold, and approves or rejects no algorithm.^{2,4} The seventh category requires quantitative measures of performance to be surfaced but does not say what they must be; the fifth requires the process used to ensure fairness in development to be described but does not require it to have worked. A module can be certified while surfacing a model with poor discrimination, provided the number is disclosed and the practices documented. The unit is also misstated: a Health IT Module is certified and a model is not, so “ONC-certified AI” is a category error.

Nor does the criterion bind clinicians or practices. A clinic does not comply with (b)(11); its EHR developer does. Providers are reached indirectly, because from January 1, 2025 the DSI criterion is what a module must hold for the CEHRT definition CMS programs rely on.¹ The rule’s title is the clue: HTI-1 is an *algorithm transparency* rule, and transparency is what it delivers.

4 The rest of the certification program

Who certifies, and who tests

Certification is conferred by accredited private bodies under federal rules, not by the agency. ONC-Authorized Certification Bodies must hold ISO/IEC 17065 accreditation; they certify Health IT Modules and conduct surveillance (§§ 170.510, 170.523); ONC-Authorized Testing Labs must hold NVLAP accreditation to ISO/IEC 17025 and perform conformance testing (§§ 170.511, 170.524). Under § 170.523(f) each ONC-ACB reports its certified modules to ASTP/ONC at least weekly, which is what populates the public Certified Health IT Product List. Two ONC-ACBs are authorized: Drummond Group and SLI Compliance.⁶

Conditions and Maintenance of Certification

Certification is not a one-time event. Seven conditions bind a certified developer continuously under §§ 170.400–170.407: information blocking (§ 170.401); assurances (§ 170.402); communications, which bars gag clauses on usability, interoperability, security, user experience, business practices and manner of use (§ 170.403); APIs (§ 170.404); real world testing (§ 170.405); semiannual attestations (§ 170.406); and the Insights Condition (§ 170.407), under which developers with 50 or more hospital sites or 500 or more clinician users report measures, phasing in during July 2027, July 2028 and July 2029.⁷

How third-party AI reaches the data

The standardized API criterion at § 170.315(g)(10) requires a certified module to expose a FHIR API for single-patient and multi-patient access, built on HL7 FHIR Release 4.0.1 with a USCDI v3 payload; HTI-1 added patient authorization revocation at § 170.315(g)(10)(vi).⁸ Paired with the API Condition at §

170.404, this is how third-party AI reaches EHR data without a bespoke integration. The consequence is structural: transparency attaches to decision support inside the certified module, while the API condition guarantees a supply line to applications outside it that carry no source-attribute obligation.

The teeth

Certification obligations run against developers; provider-side enforcement is separate. Under the information blocking disincentives final rule, 89 Fed. Reg. 54662–54718 (July 1, 2024), effective July 31, 2024, hospitals and critical access hospitals are not meaningful EHR users for the reporting period, MIPS eligible clinicians receive a zero in the Promoting Interoperability category, and accountable care organizations may be removed from the Medicare Shared Savings Program for at least a year.⁹ None is self-executing; each requires an HHS Office of Inspector General determination. Providers face disincentives; developers, exchanges and networks face civil monetary penalties.

5 The rollback now pending

This regime may not survive in its current form. HTI-5, *Health Data, Technology, and Interoperability: ASTP/ONC Deregulatory Actions To Unleash Prosperity*, 90 Fed. Reg. 60970–61034 (Dec. 29, 2025), RIN 0955-AA09, proposes removing **34 of 60** certification criteria, revising 7 and leaving 19 unchanged, and pivoting the program toward FHIR-based API requirements. The provision that matters here is explicit: it proposes to reduce the scope of the decision support interventions criterion to “fully remove the artificial intelligence (AI) ‘model card’ requirements.” Comments closed February 27, 2026.¹⁰

As of September 2026 no final rule has issued, so the source-attribute and IRM obligations remain in force. Describing them as repealed is wrong; describing them as settled law without noting the proposal is incomplete.

A separate instrument is often mistaken for relief. On November 25, 2025 ASTP/ONC announced that it would not exercise direct review authority arising solely from non-compliance with the **January 1, 2026** HTI-1 compliance date, with no enforcement before March 1, 2026. That notice covers the criteria HTI-1 revised as of January 1, 2026.¹¹ It does not reach the DSI criterion, whose date was January 1, 2025.

EVERY STATEMENT OF THESE RULES CARRIES A DATE

This area is in active flux, and the flux is not incidental to the subject: the DSI criterion is simultaneously in force and the target of a pending proposal to remove its AI provisions, and the FDA's clinical decision support guidance was reissued twice within one month in 2026. Two sourcing limits follow. The scope of the November 2025 enforcement discretion notice is described here from ASTP/ONC's published summary rather than from the notice text, and no Federal Register notice of availability was located for the January 2026 FDA guidance.

6 The FDA line

The statutory exclusion

Section 3060(a) of the 21st Century Cures Act added FD&C Act § 520(o), codified at 21 U.S.C. § 360j(o)(1)(E). A software function falls outside the device definition “unless the function is intended to acquire, process, or analyze a medical image or a signal from an in vitro diagnostic device or a pattern or signal from a signal acquisition system, for the purpose of— (i) displaying, analyzing, or printing medical information about a patient or other medical information (such as peer-reviewed clinical studies and clinical practice guidelines); (ii) supporting or providing recommendations to a health care professional about prevention, diagnosis, or treatment of a disease or condition; and (iii) enabling such health care professional to independently review the basis for such recommendations that such software presents so that it is not the intent that such health care professional rely primarily on any of such recommendations to make a clinical diagnosis or treatment decision regarding an individual patient.”³ All four elements must be satisfied.

The first is a hard gate stated in the negative. A function intended to acquire, process or analyze a medical image, a signal from an in vitro diagnostic device, or a pattern or signal from a signal acquisition system is outside the exclusion entirely, however transparent it is and however firmly a human sits in the loop. Imaging AI and electrocardiogram AI can never be non-device clinical decision support.

The fourth element does the analytical work. The software must let the professional review the basis for the recommendation independently, so that it is not intended that the professional rely primarily on the output. A confirmation button does not satisfy that. What does is a presentation of inputs, logic and sources sufficient for a competent professional to reach the same conclusion unaided, which is why opaque models are structurally difficult to fit.

Two points are regularly misstated. The exclusion applies to software *functions* rather than products, and one product routinely contains both. And enforcement discretion is not exclusion: software FDA declines to regulate actively is still a device.

The guidance history, and why the version matters

FDA's reading of those elements changed in 2026. The agency issued a revised final guidance on **January 6, 2026** and reissued it on **January 29, 2026** under Docket No. FDA-2017-D-6569, holding a public town hall on March 11, 2026.¹² The previous final guidance, notice of availability at 87 Fed. Reg. 58810 (Sept. 28, 2022), governed until January 2026 and is the version most secondary commentary still describes, including its treatment of time-critical outputs and single specific recommendations as indicators of device status.¹³

Counsel reviewing the current document report that it accepts a **single recommendation** where only one is clinically appropriate, reversing the 2022 posture; clarifies the treatment of clinical documentation and summarization tools that do not themselves analyze underlying images; emphasizes usability and well-understood and accepted sources; and moves the time-critical-decision limitation from the third criterion to the fourth without changing the outcome. That characterization comes from law-firm analyses rather than an agency comparison, and commentators criticized the guidance for leaving consumer-facing tools unaddressed.

Changing a model after authorization

Where a function is a device, the mechanism that lets its model change without a new marketing submission is the Predetermined Change Control Plan. The guidance was originally issued December 4, 2024, notice of availability at 89 Fed. Reg. 96259–96261 under Docket No. FDA-2022-D-2628; the current version is dated **August 18, 2025**. A PCCP has three parts: a **Description of Modifications**; a **Modification Protocol** covering verification and validation methods, data management, update procedures and performance monitoring; and an **Impact Assessment**.¹⁴ A genuinely non-device function needs none, and none is available.

How much of this reaches an EMR

Very little. A peer-reviewed taxonomy in *npj Digital Medicine* analyzed all **1,016** FDA authorizations of AI- and ML-enabled devices current as of December 20, 2024, covering authorizations through September 27, 2024 and resolving to **736 unique devices**. Medical imaging accounted for **84.4%** of devices and EHR-based inputs for **0.4%**, and the authors found no evidence of large language models.¹⁵ The claim that FDA already regulates the AI in an EHR does not survive those figures.

WHAT THE DEVICE FIGURES DO AND DO NOT SHOW

The 1,016 authorizations and the 0.4% EHR-input share describe devices FDA has authorized, not AI in clinical use. They support no claim about how much AI is deployed in EHRs, only about how little has passed through device review.

7 The voluntary layer and the state layer

Assurance, as proposed and as it turned out

The most-cited proposal here is a 2024 *JAMA* position paper for a nationwide network of health AI assurance laboratories: a public-private network applying community best practices to test health AI models, issuing performance reports usable across populations and settings, with lifecycle monitoring.¹⁶ It is a proposal by named authors in a journal, not a program.

What followed is reported by trade press rather than primary agency records, and is given here for direction rather than quotation. An HHS Deputy Secretary referenced a national assurance-lab network at Health Datapalooza in September 2024, but none was established and it never had a statutory or regulatory basis. The ONC and FDA officials on the Coalition for Health AI board sat as *non-voting* members and departed in July 2024; the concept was abandoned by early 2025.¹⁷ CHAI is a private 501(c)(3) with no governmental authority.

Its documents should be labeled for what they are. The *Assurance Standards Guide*, released for public comment on 26 June 2024 with a 60-day comment period, is a lifecycle framework from use-case definition through development, deployment and monitoring, illustrated with six worked use cases. It is an industry consensus document, not peer-reviewed, with no legal force, and it is not incorporated by reference into any federal rule.¹⁸ The companion *Applied Model Card*, draft version 0.1 of November 2024, cites § 170.315(b)(11) throughout and says it supports meeting HTI-1's criteria for predictive DSIs.¹⁹ The relationship runs one way: HTI-1 requires source attributes and CHAI published a template that maps to them. HTI-1 does not require a CHAI model card.

Reporting guidelines

Where regulation asks for quantitative measures of performance without saying how to report them, reporting guidelines say how. TRIPOD-LLM extends TRIPOD+AI to studies using large language models, with a modular checklist covering development, tuning, prompting and evaluation.²⁰ CONSORT-AI adds 14 items to CONSORT 2010 for trial reports and SPIRIT-AI adds 15 to SPIRIT 2013

for protocols.^{21,22} DECIDE-AI, with 27 items, covers the stage most relevant to an EMR shipping a feature: live clinical evaluation between offline validation and large comparative trials.²³ They bind authors through journal policy, and no developer through law.

States

California AB 3030 (Stats. 2024, ch. 848), operative January 1, 2025 at Cal. Health & Safety Code §§ 1339.75–1339.78, requires a health facility, clinic, physician’s office or group practice office using generative AI for patient communications pertaining to patient clinical information to include a disclaimer that the communication was AI-generated and clear instructions for contacting a human.²⁴ The exemption is the operative provision: the requirements do not apply where a licensed or certified health care provider reads and reviews the communication.

California SB 1120 (Stats. 2024, ch. 879), operative January 1, 2025, amends Health & Safety Code § 1367.01 and Insurance Code § 10123.135 and binds plans and insurers conducting utilization review, not clinicians. Determinations made with AI must rest on the enrollee’s clinical history and individual circumstances rather than a group dataset alone. A determination of medical necessity “shall be made only by a licensed physician or a licensed health care professional competent to evaluate the specific clinical issues involved,” and an AI tool “shall not deny, delay, or modify health care services based, in whole or in part, on medical necessity.”²⁵

Texas and Utah legislate disclosure and little else. Texas HB 149, the Texas Responsible Artificial Intelligence Governance Act, creates Tex. Bus. & Com. Code chapters 551–554, took effect January 1, 2026, and is far narrower than the bill as filed; its health care hook is a disclosure duty at § 552.051(f), under which a provider using AI in health care services or treatment must disclose that fact to the recipient or the recipient’s representative, clearly and conspicuously, in plain language, free of dark patterns, no later than the date the service is first provided.²⁶ Utah H.B. 452, effective May 7, 2025 at Utah Code title 13, chapter 72a, requires a mental health chatbot to disclose that it is AI and not human before a user first accesses it, at the start of any new interaction after seven days of inactivity, and whenever the user asks.²⁷

Colorado supplies the correction most needed in 2026 writing. SB 24-205, signed May 17, 2024, was the first comprehensive US AI statute, imposing duties on developers and deployers of high-risk systems making consequential decisions, health care included. SB 25B-004 pushed its effective date from February 1, 2026 to June 30, 2026, and SB 26-189, signed May 14, 2026, repealed and reenacted the act before it ever took effect. The successor regulates automated decision-making technology on a notice-and-explanation model, with key obligations beginning January 1, 2027: developer documentation on intended uses, training-data categories and known limitations; deployer notice at the point of interaction; a plain-language explanation within 30 days of an adverse decision; and a right to meaningful human review.²⁸ Any account calling Colorado’s original duties live law describes a statute that never operated.

8 What binds whom

TABLE 2

Each obligation, the party it binds, whether it is mandatory or voluntary, and the operative date. Dates are compliance or effective dates as stated in the cited instrument; status is as of September 2026.

Obligation	Binds	Force	Date
DSI criterion, § 170.315(b)(11): 13 or 31 source attributes, user access and modification	Health IT developer	Mandatory to hold certification	Certified by Dec. 31, 2024; operative Jan. 1, 2025
Intervention Risk Management and the public plain-language summary	Health IT developer, for Predictive DSIs it supplies	Mandatory	Jan. 1, 2025
Privacy and security requirements for (b)(11) modules (HTI-2)	Health IT developer	Mandatory	Jan. 1, 2028
Conditions and Maintenance of Certification, §§ 170.400–170.407	Health IT developer	Mandatory	Continuous; Insights phases July 2027, 2028, 2029
Standardized API, § 170.315(g) (10), with the API condition at § 170.404	Health IT developer	Mandatory	In force
Information blocking disincentives	Hospitals, critical access hospitals, MIPS clinicians, ACOs	Mandatory, on an OIG determination	Effective July 31, 2024
Non-device CDS exclusion, 21 U.S.C. § 360j(o)(1)(E)	Any developer of a software function	A boundary, not an obligation; failing it makes the developer a device manufacturer	In force since Dec. 13, 2016
Device requirements, including a PCCP for post-authorization model change	Device manufacturer	Mandatory once the function is a device	PCCP guidance current version Aug. 18, 2025
CHAI Assurance Standards Guide and Applied Model Card	Whoever elects to adopt them	Voluntary	Not applicable
TRIPOD-LLM, CONSORT-AI, SPIRIT-AI, DECIDE-AI	Authors and journals	Voluntary, through journal policy	Not applicable
California AB 3030 generative-AI disclaimer	Health facility, clinic, physician's office, group practice office	Mandatory, unless a licensed provider reads and reviews the communication	Jan. 1, 2025
California SB 1120 medical-necessity rule	Health plans and insurers	Mandatory	Jan. 1, 2025
Texas § 552.051(f) AI-use disclosure	Provider using AI in health care services or treatment	Mandatory	Jan. 1, 2026
Utah H.B. 452 mental health chatbot duties	Supplier of a mental health chatbot	Mandatory	May 7, 2025

Obligation	Binds	Force	Date
Colorado SB 26-189 automated decision-making duties	Developers and deployers, health care among consequential decisions	Mandatory	Key obligations Jan. 1, 2027

Two features are worth naming. Every mandatory federal obligation on a health IT developer is procedural: disclose, document, enable, publish, attest, test. And every obligation touching what a model outputs clinically lands on somebody other than the developer: a provider disclosing that AI was used, or a payer barred from letting a tool decide medical necessity.

9 Where the gap is

Transparency is regulated. Performance is not.

The only federal instrument that reaches a predictive feature inside an ambulatory EMR requires 31 attributes to be surfaced, a risk-management practice to be documented across eight named dimensions, and a plain-language summary to be posted where anyone can read it without an account. It sets no threshold for any number it requires to be disclosed. FDA device review, the regime that does examine performance, reaches an EHR-embedded predictive feature almost never.¹⁵ The layer that proposed independent evaluation was never built.^{16,17} State law adds disclosure duties for providers and a decision-authority rule for payers, and says nothing about whether a model works.

The binding requirements are therefore procedural and the substantive ones self-imposed. A developer wanting a defensible claim about how well a feature performs has to generate the evidence, choose a reporting standard nobody is enforcing, and hold a line nobody is checking. If the model-card provisions are removed as proposed, the disclosure that made an outside appraisal possible becomes optional too.

REFERENCES

Entries 1 to 14 are statutes, regulations and agency guidance rather than research: cite them for what they require, not as evidence that anything works. Entry 15 is the only peer-reviewed empirical study here. Entries 16 to 23 carry no legal force — a journal proposal, trade reporting, two voluntary industry documents and four reporting guidelines — and entry 17 is used for the direction of events, not for quotation. Entries 24 to 28 are state statutes.

1. Office of the National Coordinator for Health Information Technology. *Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing*. Final rule. 89 Fed. Reg. 1192–1438 (Jan. 9, 2024); Doc. No. 2023-28857; RIN 0955-AA03; effective Feb. 8, 2024. [federalregister.gov](https://www.federalregister.gov)

REGULATION

2. Office of the National Coordinator for Health Information Technology. *Decision support interventions*. 45 C.F.R. § 170.315(b)(11) (2026). [ecfr.gov](https://www.ecfr.gov) REGULATION
3. United States Congress. *21st Century Cures Act, § 3060(a), Clarifying Medical Software Regulation*. Pub. L. No. 114-255, 130 Stat. 1033, 1130 (Dec. 13, 2016), codified at 21 U.S.C. § 360j(o)(1)(E). [uscode.house.gov](https://www.uscode.house.gov) STATUTE
4. Assistant Secretary for Technology Policy / Office of the National Coordinator for Health Information Technology. *9 Key Functionalities and Requirements of the New (b)(11) Decision Support Interventions Certification Criterion*. DSI Explainer series, July 2024; with *HTI-1 Decision Support Interventions (DSI) Fact Sheet*, Dec. 2023. [healthit.gov](https://www.healthit.gov) GOVERNMENT REPORT
5. Office of the National Coordinator for Health Information Technology. *Health Data, Technology, and Interoperability: Trusted Exchange Framework and Common Agreement (TEFCA)*. Final rule. 89 Fed. Reg. 101772–101817 (Dec. 16, 2024); Doc. No. 2024-29163; RIN 0955-AA07; effective Jan. 15, 2025. [federalregister.gov](https://www.federalregister.gov) REGULATION
6. Office of the National Coordinator for Health Information Technology. *ONC Health IT Certification Program*. 45 C.F.R. Part 170, Subpart E, §§ 170.500–170.599. [ecfr.gov](https://www.ecfr.gov) REGULATION
7. Office of the National Coordinator for Health Information Technology. *Conditions and Maintenance of Certification Requirements for Health IT Developers*. 45 C.F.R. Part 170, Subpart D, §§ 170.400–170.407. [ecfr.gov](https://www.ecfr.gov) REGULATION
8. Office of the National Coordinator for Health Information Technology. *Standardized API for patient and population services*. 45 C.F.R. § 170.315(g)(10); standards adopted at 45 C.F.R. §§ 170.213, 170.215. [healthit.gov](https://www.healthit.gov) REGULATION
9. Centers for Medicare & Medicaid Services and Office of the National Coordinator for Health Information Technology. *21st Century Cures Act: Establishment of Disincentives for Health Care Providers That Have Committed Information Blocking*. Final rule. 89 Fed. Reg. 54662–54718 (July 1, 2024); Doc. No. 2024-13793; RIN 0955-AA05; effective July 31, 2024. [federalregister.gov](https://www.federalregister.gov) REGULATION
10. Assistant Secretary for Technology Policy / Office of the National Coordinator for Health Information Technology. *Health Data, Technology, and Interoperability: ASTP/ONC Deregulatory Actions To Unleash Prosperity*. Proposed rule. 90 Fed. Reg. 60970–61034 (Dec. 29, 2025); Doc. No. 2025-23896; RIN 0955-AA09; comments closed Feb. 27, 2026. [federalregister.gov](https://www.federalregister.gov) REGULATION
11. Assistant Secretary for Technology Policy / Office of the National Coordinator for Health Information Technology. *Certification Criteria Compliance Dates Enforcement Discretion Notice*. Announced Nov. 25, 2025. [healthit.gov](https://www.healthit.gov) GOVERNMENT REPORT

12. U.S. Food and Drug Administration. *Clinical Decision Support Software: Guidance for Industry and Food and Drug Administration Staff*. Final guidance issued Jan. 29, 2026, superseding the version issued Jan. 6, 2026; Docket No. FDA-2017-D-6569. [fda.gov](https://www.fda.gov) GUIDANCE
13. U.S. Food and Drug Administration. *Clinical Decision Support Software; Guidance for Industry and Food and Drug Administration Staff; Availability*. 87 Fed. Reg. 58810 (Sept. 28, 2022); Doc. No. 2022-20993; Docket No. FDA-2017-D-6569. [federalregister.gov](https://www.federalregister.gov) GUIDANCE
14. U.S. Food and Drug Administration. *Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence-Enabled Device Software Functions*. Final guidance originally issued Dec. 4, 2024; current version Aug. 18, 2025. Notice of availability: 89 Fed. Reg. 96259–96261 (Dec. 4, 2024); Docket No. FDA-2022-D-2628. [fda.gov](https://www.fda.gov) GUIDANCE
15. Singh R, Bapna M, Diab AR, et al. *How AI is used in FDA-authorized medical devices: a taxonomy across 1,016 authorizations*. npj Digital Medicine. 2025;8:388. [doi:10.1038/s41746-025-01800-1](https://doi.org/10.1038/s41746-025-01800-1) CROSS-SECTIONAL
16. Shah NH, Halamka JD, Saria S, et al. *A Nationwide Network of Health AI Assurance Laboratories*. JAMA. 2024;331(3):245–249. [doi:10.1001/jama.2023.26930](https://doi.org/10.1001/jama.2023.26930) POSITION PAPER
17. Fierce Healthcare. *Inside CHAI's failed assurance labs*. Trade reporting on the Coalition for Health AI, its federal relationships and the abandonment of the assurance-lab concept. [fiercehealthcare.com](https://www.fiercehealthcare.com) JOURNALISM
18. Coalition for Health AI. *Assurance Standards Guide* and companion *Assurance Reporting Checklists*. Released for public comment 26 June 2024, with a 60-day comment period. chai.org STANDARD
19. Coalition for Health AI. *The CHAI Applied Model Card*. Draft v0.1, November 2024; comment window 22 Nov. 2024 to 22 Jan. 2025. mc.chai.org STANDARD
20. Gallifant J, Afshar M, Ameen S, et al. *The TRIPOD-LLM reporting guideline for studies using large language models*. Nature Medicine. 2025;31(1):60–69. [doi:10.1038/s41591-024-03425-5](https://doi.org/10.1038/s41591-024-03425-5) REPORTING GUIDELINE
21. Liu X, Cruz Rivera S, Moher D, et al. *Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension*. Nature Medicine. 2020;26(9):1364–1374. [doi:10.1038/s41591-020-1034-x](https://doi.org/10.1038/s41591-020-1034-x) REPORTING GUIDELINE
22. Cruz Rivera S, Liu X, Chan A-W, et al. *Guidelines for clinical trial protocols for interventions involving artificial intelligence: the SPIRIT-AI extension*. Nature Medicine. 2020;26(9):1351–1363. [doi:10.1038/s41591-020-1037-7](https://doi.org/10.1038/s41591-020-1037-7) REPORTING GUIDELINE
23. Vasey B, Nagendran M, Campbell B, et al. *Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI*. Nature Medicine. 2022;28(5):924–933. [doi:10.1038/s41591-022-01772-9](https://doi.org/10.1038/s41591-022-01772-9) REPORTING GUIDELINE

24. California Legislature. *Assembly Bill No. 3030, Health care services: artificial intelligence*. Stats. 2024, ch. 848; filed 28 Sept. 2024; operative 1 Jan. 2025; codified at Cal. Health & Safety Code §§ 1339.75–1339.78. [leginfo.legislature.ca.gov](https://leginfo.ca.gov) STATUTE
25. California Legislature. *Senate Bill No. 1120, Health care coverage: utilization review*. Stats. 2024, ch. 879; approved 28 Sept. 2024; operative 1 Jan. 2025; amending Cal. Health & Safety Code § 1367.01 and Cal. Insurance Code § 10123.135. [leginfo.legislature.ca.gov](https://leginfo.ca.gov) STATUTE
26. Texas Legislature. *Texas Responsible Artificial Intelligence Governance Act*. H.B. 149, 89th Leg., R.S. (2025); Tex. Bus. & Com. Code tit. 11, subtit. D, chs. 551–554; effective 1 Jan. 2026. capitol.texas.gov STATUTE
27. Utah Legislature. *Artificial Intelligence Amendments*. H.B. 452 (2025 Gen. Sess.); enacting Utah Code tit. 13, ch. 72a and § 58–60–118; effective 7 May 2025. [le.utah.gov](https://leg.utah.gov) STATUTE
28. Colorado General Assembly. *Consumer Protections for Artificial Intelligence* (SB 24–205, signed 17 May 2024); *Increase Transparency for Algorithmic Systems* (SB 25B–004, signed 28 Aug. 2025); *Automated Decision-Making Technology* (SB 26–189, signed 14 May 2026, repealing and reenacting the Colorado AI Act, key obligations 1 Jan. 2027). leg.colorado.gov STATUTE

RELATED PAPERS

Approval gates

A specification for where human decision authority has to sit.

Human oversight of clinical AI

Whether the reviewer the rules rely on can actually catch the error.

Automation in the revenue cycle

Coding, denials and prior authorization, and the rules now landing on payers.