

FDA Wants Hospitals to Help Police Generative AI Devices After Launch

Federal regulators propose a doctor style competency exam for generative AI devices, shifting postmarket safety monitoring onto hospital clinical teams.

SERIES Q-Tips · READ 5 min · STATUS Board approved

Medical device regulation is the federal rule testing clinical software before it reaches a hospital floor. Under traditional rules, the Food and Drug Administration reviews a vendor filing, confirms fixed inputs produce reliable outputs, and clears the product.

Originally, this review was static. Engineers fed charts into a deterministic rules tool where input A always equals output B, proving accuracy in advance so flawed code stayed out of exam rooms.

Generative artificial intelligence broke that limit. Unlike a deterministic rules tool, a generative model runs an auto regressive loop inside a context window, predicting the next word from probabilities and producing varied responses that can hallucinate errors. Because input and output testing cannot cover a generative model, the agency had not authorized any generative AI based device for any clinical purpose as of its November 6th, 2025 Digital Health Advisory Committee meeting ([Hogan Lovells](#), December 2025).

On August 18th, 2026, the agency released a discussion paper proposing a competency based model (device benchmarking plus clinical confirmation) and shared responsibility for postmarket oversight by health care institutions, clinicians, and payers, with comments on docket [FDA-2026-N-7874](#) due October 19th, 2026 ([Food and Drug Administration](#), August 2026; [Cooley](#), September 2026). Commenters on the Department of Health and Human Services request for information asked above all for evaluation and benchmarking tools and governance guidance ([Healthcare Dive](#), June 2026). Instead, regulators say sample based clinician review and drift monitoring after deployment are the only workable control, shifting the regulator burden onto unpaid clinical adjudicators with no credentialing, oversight, or liability shield ([Cooley](#), September 2026). We ran the issue through Contextualism, our framework for training AI to grasp the full picture of a complex issue by stress testing it across four quadrants simultaneously.

The four filters

PROTAGONIST: FOR COMPETENCY BASED REGULATION. The protagonist argues for the proposal: generative software acts like a trainee, so regulators must evaluate it like a doctor. Pairing device benchmarking with clinical confirmation clears helpful tools while relying on hospitals to catch local drift early ([Cooley](#), September 2026).

ANTAGONIST: AGAINST SHIFTING POLICING ONTO HOSPITALS. The antagonist argues against offloading oversight: federal officials are shifting the regulator burden onto unpaid clinical adjudicators.

Clinicians must grade real world samples against preset criteria, yet the paper offers no credentialing, oversight, or liability shield ([Cooley](#), September 2026).

WITNESS: THE PATIENT AND CLINICAL REALITY. The witness presents everyday reality: generative software already writes charts while rules lag. After ambient documentation drafts a visit note, a resident signs 20 discharge summaries at midnight and misses one fabricated dose while patients read portal replies without knowing if a doctor wrote the advice.

UTILITARIAN: THE MACROECONOMIC AND SYSTEMIC VIEW. The utilitarian evaluates the net outcome: forcing thousands of hospitals to run separate postmarket audits wastes clinical capacity. Society gains when regulators mandate vendor funded surveillance, shared Foundation Model Master Files, and national drift reporting so one safety signal protects 200 health systems at once ([Cooley](#), September 2026).

Where the truth lives

Both sides agree input and output testing cannot cover a generative model in local electronic health records. They split over who funds, staffs, and bears liability for postmarket monitoring and oversight. Zero generative AI devices held clinical authorization as of the November 6th, 2025 Digital Health Advisory Committee meeting ([Hogan Lovells](#), December 2025), while commenters on the Department of Health and Human Services request for information asked above all for evaluation and benchmarking tools and governance guidance ([Healthcare Dive](#), June 2026).

How AI helps with this

Leaders can use the general assistant your organization licenses to build their response before October 19th, 2026. Unlike a deterministic rules tool, a generative model runs an auto regressive loop inside a context window and can hallucinate errors. Under a Business Associate Agreement (BAA), a team can analyze vendor contracts and stripped outputs to inventory every generative AI function in use and place each on the agency two axis grid (degree of autonomy against severity of a wrong output). From that inventory, the tool drafts the comment letter and a sample based clinician review protocol for one tool, which the chief medical information officer and general counsel review and sign while keeping patient information examples out.

The solution

Health systems need a four part defense governed by **The Name Standard**, requiring a licensed leader with time, authority, and refusal rights to review and sign every AI assisted filing:

- **Two axis grid inventory:** Place every generative AI function across ambient documentation, inbox messaging, discharge summaries, and revenue cycle on the grid by autonomy against severity of a wrong output.
- **Docket comment letter:** Draft the letter on docket FDA-2026-N-7874 from that inventory before October 19th, 2026, demanding credentialing, oversight, and liability shields.
- **Sample based clinician review protocol:** Pick one tool, strip all 18 identifiers from weekly outputs under the Safe Harbor Method, and have two clinicians score the sample for drift.
- **Vendor contract cost shifting:** Require vendors under a Business Associate Agreement (BAA) to share Foundation Model Master Files, device benchmarking scores, and postmarket review funding.

Operational moves for health system leaders

- 1. Inventory every generative AI function in 48 hours.** Have your chief medical information officer place every live function on the two axis grid by autonomy against severity of a wrong output.
- 2. Analyze vendor contracts and outputs without patient data.** Feed vendor contracts and stripped outputs, never patient information examples, into your licensed AI tool to expose unfunded duties.
- 3. Review and sign your docket letter by October 19th, 2026.** Have your chief medical information officer and general counsel review and sign your comment letter under The Name Standard.
- 4. Run a sample based clinician review protocol on one tool.** Audit 20 stripped outputs weekly on one tool so your committee holds drift data before renewal.

The decision for leadership

- 1. Stay silent while regulators shift policing onto your clinicians.** Let the October 19th, 2026 docket close without input, and absorb unpaid physician hours auditing vendor code.
- 2. Put your risk grid on the record and lock vendor contracts.** Place every generative function on the two axis grid, file your signed letter, and make vendors fund postmarket review.

The executive takeaway

Federal regulators want to credential generative AI like a doctor by pairing premarket competency testing with postmarket clinician review inside your hospital ([Cooley](#), September 2026). Place every generative AI function you run on the agency two axis grid this week, have your chief medical information officer and general counsel sign your docket letter before October 19th, 2026, and make vendors fund postmarket oversight.

Sources

- Food and Drug Administration, [Considerations for the Regulation of Generative AI Enabled Medical Devices](#), August 2026
- Regulations.gov, [Docket FDA-2026-N-7874](#), August 2026
- Cooley LLP, [Regulating AI Like a Doctor](#), September 2026
- Hogan Lovells, [FDA Digital Health Advisory Committee Weighs Guardrails for Generative AI](#), December 2025
- Healthcare Dive, [HHS Details Takeaways from Sweeping AI Request for Information](#), June 2026

How this brief was built

We ran this issue through the four filters of Contextualism: Protagonist, Antagonist, Witness, and Utilitarian. We weighed the room against our eight human values hierarchy, led by life over death, followed by health over sickness, peace over war, freedom over slavery, prosperity over poverty, knowledge over ignorance, cooperation over competition, and fairness over bias. Completed October 13th, 2026 for publication October 15th, 2026 at high confidence, this conclusion would change if final rules fund third party adjudication or shield hospitals from liability.

Do this at your organization

Members can access Place every generative AI tool you run on FDA's risk grid before the docket closes on the Q-Kit shelf in the QuadCo app. Email us the word [DOCKET] at info@thequadco.ai and the kit is in your inbox.