

MEDICAL DEVICES · FDA GENERATIVE AI

Agentic AI Is Entering Medical *Device Regulation*

FDA's new generative AI framework brings runtime AI governance and autonomous medical devices onto the regulatory and policy agenda.

MAUREEN DOYLE-SPARE

© 2026 Maureen Doyle-Spare.

SWIPE TO READ →

Agentic AI Is Entering Medical *Device Regulation*

FDA's new generative AI framework brings runtime AI governance and autonomous medical devices onto the regulatory and policy agenda.

MAUREEN DOYLE-SPARE

Agentic AI adds a new question to medical-device assurance. Not only whether the system works, but whether it acted on the meaning the institution actually authorized.

FDA has opened the question

On August 18, FDA issued a discussion paper[1] seeking public input on generative AI-enabled medical devices, covering risk assessment, premarket evaluation, post-market monitoring and other emerging issues. The paper's fourth topic area is given over to two subjects in particular: foundation models and agentic AI systems. Comments are due October 19, 2026.

FDA is careful about what the paper represents. It is exploratory, not draft or final guidance, and it does not establish new regulatory expectations. The agency is asking stakeholders to help inform its thinking as generative AI introduces risks and capabilities that differ from traditional software and existing AI-enabled medical devices.

That is what makes the discussion consequential. As medical devices move from producing predictions or recommendations toward systems capable of interpreting conditions, coordinating tasks and taking actions, assurance has to account for more than whether the underlying model is competent.

Competence and authority are different things

FDA is exploring how GenAI-enabled devices should be evaluated before market and monitored after deployment. Those are necessary questions. Agentic systems add another control problem, because a system can perform well in testing and still resolve the operational meaning of a specific case differently from the meaning the regulated organization intended.

Consider clinical eligibility. An agent may consult an approved protocol, an amendment, site guidance and patient data before determining whether a patient is eligible. Every source may be valid. The system may function exactly as designed. The interpretation it synthesizes across them can still fall outside what the sponsor authorized.

This is the boundary between traditional validation and reasoning-layer governance identified in my Life Sciences research[2]. Validation governs whether software

© 2026 Maureen Doyle-Spare.

performs as intended. Runtime governance addresses the interpretation the system forms and acts on during execution. Agentic AI makes that distinction increasingly consequential, because interpretation is moving into the operational workflow rather than remaining with a qualified human reviewer.

Validation governs whether software performs as intended. Runtime governance addresses the interpretation the system forms and acts on during execution.

What postmarket monitoring cannot currently see

FDA is also seeking input on postmarket monitoring. Existing approaches can identify performance degradation, adverse events, software failures and changes in model behavior. Agentic systems introduce a form of change that is harder to observe: the technology can remain stable while the operational meaning it resolves begins to shift.

That is the condition I have described as Agentic Workflow Drift. An autonomous system may repeatedly reconcile individually valid but semantically inconsistent sources until its working interpretation of eligible, serious, expected or critical no longer reflects the definition the institution approved.

The relevant evidence therefore extends beyond the final output. For consequential decisions, an institution needs to preserve what the agent actually resolved at runtime and compare it against the definition it was authorized to use.

The Semantic Control Plane[3] is one architectural approach to that problem. The institution holds an authorized definition for the decision. At runtime, the system captures the interpretation the agent has actually formed, measures how far it has diverged from the authorized version, and applies the institution's threshold before the action is permitted to proceed. The record of that evaluation is preserved so the decision can be reconstructed later. This is not an argument that FDA should adopt the SCP as a regulatory framework. It illustrates the kind of runtime evidence that becomes relevant as autonomous systems begin making consequential interpretations during operation.

Autonomy raises the stakes at the execution boundary

A generative medical device that drafts text for a clinician occupies a different control position from an agentic system that sequences actions, invokes tools, prioritizes cases or drives downstream workflows.

As autonomy increases, the boundary between reasoning and execution matters more. Measurement alone does not create governance. An institution needs an enforceable decision about whether a particular interpretation is authorized to proceed, made before the action is taken rather than reconstructed afterward.

© 2026 Maureen Doyle-Spare.

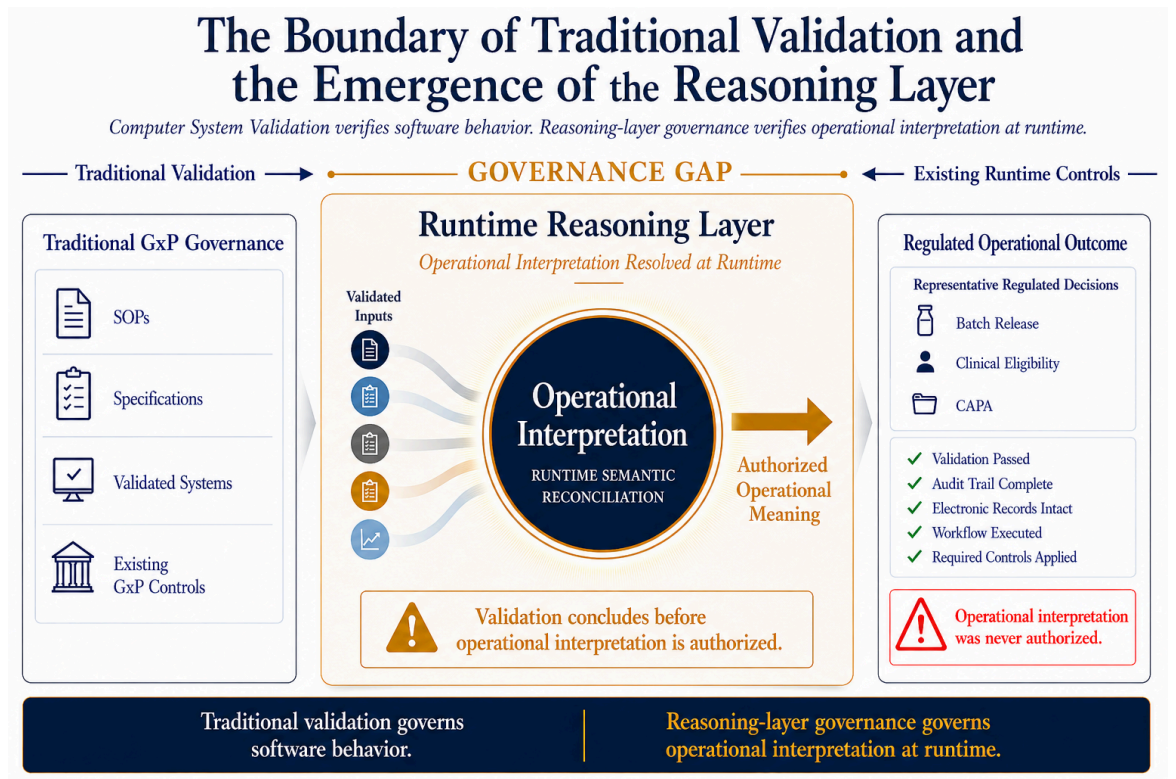


Figure 1. *The Boundary of Traditional Validation and the Emergence of the Reasoning Layer. Source: Doyle-Spare (2026)* [2]

The principle extends beyond medical devices

FDA's discussion is specific to medical devices, but the underlying governance issue is cross-industry. A bank may need to govern how an autonomous system resolves approved or within tolerance, an insurer how it resolves coverage, a manufacturer how it resolves critical or release, a critical-infrastructure operator whether an interpretation of an operating condition carries authority to proceed.

In each case, stronger model evaluation remains necessary. It simply governs a different object. The same separation applies to cybersecurity. If an adversary deliberately targets the semantic conditions under which an agent forms its interpretation, the exposure can enter the Semantic Layer Integrity Attack[4] threat surface. Conventional cybersecurity controls remain essential, but they do not remove the need to govern the meaning under which an otherwise authorized system acts.

The comment window is open

FDA is deliberately seeking stakeholder input before establishing regulatory expectations, and is asking device manufacturers, clinicians, researchers, the public and other interested parties to contribute. Comments are due October 19, 2026, under docket FDA-2026-N-7874[5] on Regulations.gov.

That creates a legitimate opportunity to put reasoning-layer governance into the record without overstating what FDA has concluded.

The contribution is narrow. Existing medical-device regulation already includes substantial premarket, lifecycle, quality, cybersecurity and postmarket controls. Agentic AI adds a distinct runtime question: when an

© 2026 Maureen Doyle-Spare.



Figure 2. *The Pre-Execution Assurance Protocol. Source: Doyle-Spare (2026)*[3]

autonomous system forms an interpretation and uses it to drive a consequential action, what evidence establishes that the meaning itself was authorized?

Generative AI is changing how regulators think about competence and performance. Agentic AI brings authority into that discussion. As autonomy increases, governing what a system can do will need to sit alongside governing the meaning under which it is allowed to act. ◆

References

1 *U.S. Food and Drug Administration. "Considerations for the Regulation of Generative AI-Enabled Medical Devices: Discussion Paper and Request for Feedback." August 18, 2026. fda.gov/medical-devices/digital-health-center-excellence/considerations-regulation-generative-ai-enabled-medical-devices-discussion-paper-and-request*

2 *Doyle-Spare, Maureen. Agentic Workflow Drift in Life Sciences: Extending the Reasoning-Layer Risk Taxonomy to GxP-Regulated Pharmaceutical and Biotechnology Operations. SSRN, July 6, 2026. doi.org/10.2139/ssrn.7301980*

3 *Doyle-Spare, Maureen. Agentic AI Systems Governance: A Runtime Reference Architecture for the Reasoning Layer and the Semantic Control Plane in Regulated Financial Institutions. Zenodo, 2026. doi.org/10.5281/zenodo.20749051*

4 *Doyle-Spare, Maureen. Agentic AI Cyber Subversion: The Semantic Layer Integrity Attack as a New Threat Class Against the Reasoning Layer. SSRN No. 6926219. doi.org/10.2139/ssrn.6926219*

5 *U.S. Food and Drug Administration. Docket No. FDA-2026-N-7874, Regulations.gov. regulations.gov/docket/FDA-2026-N-7874*

Originally published in *The Agentic AI Governance Playbook* on Substack, August 20, 2026. maureendoylespare.substack.com

© 2026 Maureen Doyle-Spare.



As autonomy increases, governing what a system can do will need to sit alongside governing the meaning under which it is allowed to act.

THE AGENTIC AI GOVERNANCE PLAYBOOK

Download the research

Papers, briefs and reference architecture at maureendoylespare.com

Visit the Playbook

WEEKLY NEWSLETTER

Follow the newsletter

The Agentic AI Governance Playbook Weekly Newsletter

Follow Newsletter

INTELLECTUAL PROPERTY: © 2026 Maureen Doyle-Spare. All rights reserved. No part of this article, including its text, figures, tables, and images, may be reproduced or distributed without the author's prior written permission.