



October 19, 2026

Dr. Michelle Tarver
Director, Center for Devices and Radiological Health
U.S. Food and Drug Administration
10903 New Hampshire Avenue
Silver Spring, MD 20993

Re: Docket No. FDA-2026-N-7874 – Considerations for the Regulation of Generative AI-Enabled Medical Devices: Discussion Paper and Request for Feedback

Dear Dr. Tarver:

The U.S. Chamber of Commerce ("the Chamber") appreciates the opportunity to provide feedback on the Food and Drug Administration's ("FDA") Discussion Paper on "Considerations for the Regulation of Generative AI-Enabled Medical Devices" (the "Discussion Paper"). The Chamber is the world's largest business federation, representing businesses of all sizes, sectors, and regions, including medical device manufacturers, health technology developers, foundation model providers, healthcare systems, and employers that provide health coverage to millions of American workers and their families.

Generative AI-enabled medical devices hold real promise for patients and the health system — improving diagnostic accuracy, expanding access to care in underserved communities, and reducing the administrative burden on clinicians. The Chamber supports FDA's goal of ensuring these technologies are safe and effective. But the regulatory framework FDA develops here will shape whether American companies can invest in and bring these products to market. A framework that is poorly calibrated, in either direction, risks delaying patient access to beneficial tools, chilling investment, or ceding U.S. leadership in medical AI to competitors operating under more appropriate regulatory regimes. Equally important, the framework should reflect the experience of healthcare providers implementing these technologies in complex clinical workflows, with patient safety, clinician oversight, and operational feasibility central to regulatory design.

The Chamber does not respond to every question in the Discussion Paper, but focuses its comments on those areas of greatest consequence to our members. These comments draw on feedback from a range of stakeholders across the healthcare field, including medical device manufacturers, health technology developers, foundation model providers, and healthcare systems. They also reflect discussions from a July 2026 Chamber roundtable on AI in healthcare, which brought together more than



fifty leaders from government, industry, and the healthcare sector, including representatives from FDA, HHS, CMS, and ONC. That discussion highlighted the need for stronger public-private coordination, greater regulatory clarity, and sustained trust-building as AI is deployed across the healthcare system — themes reflected throughout these comments. We look forward to continued engagement with CDRH as this framework develops.

I. The Two-Axis Framework Should Be Refined and Used to Inform Risk Management and Evidence Expectations (Questions 1, 2, 3, and 4)

The Chamber supports organizing risk assessment around a framework that considers both how independently a device acts and how consequential an incorrect output would be. That is an appropriate conceptual approach for a technology category as diverse as generative AI-enabled medical devices, which range from patient-facing informational tools to systems that can initiate clinical orders. The two-axis structure is a meaningful conceptual advance. We also note that the use of "device" should not include tools or functions that do not meet the statutory definition of a medical device, including certain software functions.

FDA should, however, distinguish clearly among the different regulatory functions that may be informed by a risk framework. Factors such as the potential consequences of an incorrect output, the degree of user reliance, and the availability of meaningful human oversight are relevant to hazard identification, risk controls, validation, residual-risk assessment, postmarket monitoring, and the type and amount of evidence necessary to demonstrate reasonable assurance of safety and effectiveness.

These considerations do not, by themselves, determine a device's statutory classification, the availability of a predicate, or the appropriate premarket pathway. Those determinations remain governed by established statutory and regulatory principles. Accordingly, the two-axis framework would be most useful as a tool to inform risk management and evidence-generation expectations, rather than as a mechanism for assigning regulatory consequences that are otherwise determined through established classification, substantial-equivalence, and premarket review frameworks.

FDA should avoid allowing a high-level risk characterization to operate as a de facto new classification system or as an automatic trigger for a particular regulatory pathway.

As currently described, however, the framework does not give manufacturers or FDA reviewers a consistent basis for placing a device on either axis. The line between "non-directive" and "action-directing" is not self-evident, and the Discussion Paper's acknowledgment that a disclaimer like "talk to your doctor" does not make an action-directing output less directive — while correct — underscores the need for affirmative criteria, not just negative ones.



The Chamber urges CDRH to develop clear, objective criteria for axis placement, including worked examples across clinical domains, in collaboration with industry, clinical stakeholders, and relevant standards-development organizations. Those criteria should account for the potential consequences of both over-escalation—directing patients to seek care when it is not needed—and under-escalation—failing to direct patients to seek care when it is needed. Both can create patient-safety and health-system consequences, and the risk assessment should reflect the potential harms in each direction. Without that clarity, manufacturers may characterize similar devices differently, and FDA reviewers may lack a common standard against which to evaluate submissions. Before applying this risk framework, FDA should first determine whether each software function meets the statutory definition of a device, including the exclusions in section 520(o) of the Federal Food, Drug, and Cosmetic Act. The Chamber encourages FDA to clarify that the use of generative AI, the clinical setting, or the presence of implementation safeguards should not, by themselves, be treated as dispositive in determining whether a function meets the statutory definition of a device.

The Chamber is also concerned that a grid-based framework, without careful attention to how requirements scale across it, will create sharp discontinuities — situations where a small difference in how a device is characterized leads to dramatically different regulatory burdens. That creates perverse incentives: manufacturers may design around category boundaries rather than around actual risk.

Regulatory expectations should scale proportionately with a device's risk profile and should be tied to mitigations, validation, performance evidence, and postmarket controls appropriate for the intended use. FDA should avoid abrupt category lines that produce disproportionate regulatory consequences without a commensurate difference in actual risk.

Finally, the Chamber urges CDRH not to treat patient-facing functions as presumptively higher risk simply because the end user is a patient. What matters is whether the specific output, in its intended clinical context, is subject to adequate safeguards before it influences a clinical decision. A patient-facing tool that surfaces information for discussion at a future clinical visit presents a materially different risk profile than one that recommends an immediate medication adjustment. The consequences axis should reflect the actual pathway from output to clinical action, including the presence and qualifications of any human reviewer in that pathway as well as clinical urgency.

II. Evaluation Standards Should Be Developed in Robust Collaboration With Industry and Standards Bodies, Not by FDA Alone (Questions 7, 8, 9, 10, 11, 12 and 16)

The Discussion Paper proposes evaluating generative AI-enabled devices through a framework combining non-clinical benchmarking with clinical confirmation. The Chamber supports the development of rigorous, scalable, and scientifically sound evaluation methods. For generative AI-



enabled devices, however, evaluation should focus on demonstrated performance, reliability, limitations, and safety under defined intended-use conditions rather than on broad claims of generalized "competency". FDA should provide illustrative examples of when shadow-mode evaluation may generate sufficiently reliable evidence of performance, clinical workflow fit, and risk to support review of an AI-enabled device.

FDA should instead encourage task-specific, intended-use performance evaluation that identifies the clinical tasks a device is expected to perform safely, such as gathering relevant patient information, developing an assessment, identifying appropriate treatment options, recognizing when escalation is warranted, or responding appropriately as new information emerges. Evaluation should identify where a system performs reliably, where it does not, what uncertainties remain, and what mitigations are necessary. This approach would enable meaningful performance assessment without implying that a system has human-like or generalized competence.

A competency-based assessment may also assume that a suitable benchmark or comparator exists. For novel capabilities, however, there may be no validated benchmark or established comparator against which to measure performance, including where a generative AI-enabled device does not measure the same construct in the same manner as an existing clinical measure. CDRH should share its current thinking, with illustrative examples, on the evidence that may be sufficient when no established benchmark exists, including when alternative comparators, such as structured clinician-panel assessment, may be appropriate.

FDA should also explain how it intends to evaluate and verify confabulations or hallucinations in any competency-based assessment. Because inaccurate outputs may be presented persuasively and could mislead patients, healthcare providers, or evaluators, FDA should clarify the verification methods, reviewer expertise, and other safeguards it expects to be used in assessing such outputs.

The Chamber is concerned, however, about who develops the relevant benchmarks. Standards written by FDA alone, without robust input from outside experts and stakeholders, may not keep pace with a field in which the underlying technology evolves rapidly. Such standards also may not incorporate the full range of clinical, technical, operational, and implementation expertise residing in industry, academia, healthcare delivery organizations, and the clinical community.

The Chamber urges CDRH to formally partner with established multi-stakeholder standards bodies — including AAMI, IEEE, HL7, and ISO — to develop and maintain the benchmark library, rather than developing evaluation criteria unilaterally.



FDA should also provide a transparent process for recognizing consensus standards and for updating recognized evaluation approaches as technologies, clinical practice, and evidence evolve.

The Discussion Paper appropriately recognizes that clinical confirmation need not require a prospective clinical study in every case. However, without clear guidance specifying when lower-burden approaches—such as retrospective evaluation, real-world evidence review, simulation, usability testing, or structured expert assessment—may be sufficient, manufacturers may default to the most burdensome option merely to avoid regulatory uncertainty.

The Chamber urges CDRH to specify, through guidance and worked examples, the types of clinical and non-clinical evidence that are generally appropriate for different intended uses and risk profiles. These expectations should be sufficiently clear to provide predictability while preserving FDA's ability to request additional evidence when a specific and documented safety concern warrants it.

For devices designed to support clinical decision-making rather than replace it, the evaluation framework must reflect how the device is actually used. Testing a clinical decision support tool as if it operates in isolation, without accounting for the clinician's role in reviewing, accepting, or rejecting its output, does not accurately represent the risk profile of the device as deployed. The Chamber urges CDRH to develop evaluation protocols that assess device performance in the context of the full clinical workflow, including the qualifications of the intended user and the safeguards built into the care pathway.

If CDRH determines that independent third-party testing is appropriate for certain device categories, any such program must be voluntary and must include appropriate confidentiality protections, as further discussed below. Further, any such program should be implemented in accordance with established FDA processes such as Accreditation Scheme for Conformity Assessment (ASCA).

III. For Lower-Risk Devices, Robust Postmarket Monitoring Should Justify a Lighter Premarket Burden (Questions 13, 14, 15, 17, and 18)

Question 18 raises the general question whether the evidence burden for generative AI-enabled devices should shift from premarket to postmarket monitoring. This question is highly consequential for manufacturers and investors. In the Chamber's view, for lower-risk devices, a demonstrated commitment to robust postmarket monitoring should be able to justify a materially lighter premarket burden, consistent with FDA's existing statutory standards, and FDA should provide clear guidance on when and how this flexibility would apply. This flexibility should supplement, not replace, risk-appropriate premarket clinical validation, especially for tools that materially influence diagnosis or treatment. Evidence expectations should account for representative patient populations and clinical settings.



The case for this trade-off is straightforward. Generative AI-enabled devices produce variable outputs that cannot be fully characterized through premarket testing alone. The most meaningful evidence of real-world performance — how a device behaves across the full diversity of patients, clinical settings, and use cases it will encounter in practice — can only be generated postmarket. As to lower-risk and moderate-risk devices in particular, requiring extensive premarket evidence for a technology whose real-world performance is inherently difficult to predict premarket creates a serious likelihood of imposing costs without commensurate safety benefit, while delaying patient access to potentially beneficial tools.

The Chamber recommends that CDRH develop a pathway through which manufacturers of lower- and moderate-risk devices may elect a streamlined premarket review based on appropriate task-specific performance testing, validation of risk controls, and a binding postmarket monitoring plan. FDA already has multiple tools for postmarket oversight and surveillance. To the extent appropriate for a device's intended use and risk profile, FDA should leverage those authorities to enable a predictable, monitoring-forward approach rather than creating duplicative or unnecessarily novel oversight mechanisms. A qualifying monitoring plan should specify performance thresholds, monitoring frequency, reporting obligations, and the conditions under which the manufacturer will modify, suspend, or withdraw the device. FDA should specify in advance what a qualifying monitoring plan must include, so that manufacturers can design their monitoring infrastructure before market entry rather than negotiating its terms after the fact.

For higher-risk devices, the Chamber does not oppose more extensive premarket review, but urges FDA to specify the evidence standards for these devices with the same clarity it applies to lower-risk pathways. Manufacturers of higher-risk generative AI devices need to know, before they invest in clinical development, what evidence FDA will require and how it will be evaluated.

IV. Change Control Rules Must Reflect How These Products Are Actually Built and Maintained (Questions 19, 20, 21, 22, 23, and 24)

Generative AI-enabled medical devices may be updated frequently through changes to model weights, system prompts, prompt templates, prompt-routing logic, output filters, retrieval sources, and other components. The existing Predetermined Change Control Plan framework provides an important starting point. FDA should fully implement that framework consistent with its statutory authority and ensure that it can be used in a practical manner for generative AI-enabled devices. In particular, FDA should clarify how changes to system prompts, prompt templates, prompt-routing logic, and other prompt-based controls should be addressed under a PCCP. These changes may appear operationally minor but can materially affect a device's outputs, clinical behavior, performance, and risk profile. FDA guidance should distinguish between prompt changes that are appropriately managed within an approved PCCP and prompt changes that may require a new premarket submission.



Manufacturers should be able to plan ahead by submitting a change-control plan that describes the categories of anticipated updates, the boundaries of those updates, and the internal evaluation protocols that will apply. Updates that fall within the scope of an approved plan and that are supported by appropriate validation—which may include certain fine-tuning, output-filter adjustments, retrieval-source updates, or limited prompt changes, depending on the scope of the approved plan—generally should not require a new premarket submission. FDA guidance should also recognize the distinct roles of foundation-model developers, device manufacturers, and healthcare organizations in the deployment of AI-enabled medical devices, and encourage clear, appropriately enforceable arrangements for information sharing, change communication, implementation, and risk management based on the functions each party controls, consistent with existing legal requirements regarding device manufacturers' responsibilities concerning the safety and effectiveness of finished devices.

Guidance should make clear that "prompt optimization" is not categorically minor and should be evaluated according to its potential effect on the finished device. The Chamber urges CDRH to publish clear criteria for what constitutes a significant change that triggers a new submission, specific enough to give manufacturers meaningful advance certainty. Those criteria should expressly address changes to prompts and prompt-management systems, including where prompts affect clinical recommendations, output restrictions, escalation pathways, human-review mechanisms, or the device's interaction with external tools and data sources. Approved change plans should also require timely communication to deploying healthcare organizations for clinically material changes, updated version and performance documentation, and sufficient information for local workflow validation and safe implementation.

A related and more difficult issue involves updates made by third-party foundation model developers. Device manufacturers may have limited ability to prevent, delay, or fully understand updates made by third-party foundation-model developers, yet they remain responsible for ensuring that the finished medical device continues to meet applicable requirements for safety and effectiveness.

FDA should provide clear, risk-based guidance regarding when changes to a third-party foundation model may warrant additional evaluation, FDA engagement, or a new premarket submission for the finished medical device. Manufacturers should retain flexibility to manage upstream changes through appropriate technical, quality-system, and, where feasible, contractual controls, based on the device's intended use, risk profile, and the manufacturer's ability to control or obtain information about the underlying model.

Manufacturers should maintain risk-based processes to monitor the finished device, assess whether a foundation-model or other upstream change materially affects device performance or risk, conduct re-benchmarking, validation, or other evaluation proportionate to the nature and potential impact of a change, and implement appropriate corrective actions when warranted. Where an upstream change



results in a material change in device performance or risk, manufacturers should promptly notify affected healthcare providers, maintain clear version documentation, make relevant revalidation results available as appropriate, and implement corrective actions proportionate to the identified impact. FDA should recognize that a manufacturer's ability to obtain advance notice, control model versioning, or identify a specific upstream change may vary materially across deployment models. These practical constraints should inform the expectations FDA applies, while device manufacturers remain responsible for ensuring that the finished device continues to meet applicable requirements for safety and effectiveness.

FDA should also recognize that the appropriate response may vary depending on the manufacturer's technical ability to control the model version, the significance of the change, the device's intended use, the availability of mitigations, and the potential clinical consequences. Clear expectations in this area would enable manufacturers to establish appropriate technical, quality-system, and, where feasible, contractual controls before deployment.

The Chamber encourages CDRH to provide additional clarity regarding the scope and limits of its regulatory authority as applied to foundation-model developers who are not themselves engaged in manufacturing or distributing a finished medical device. Additional clarity on this topic would reduce uncertainty that may otherwise discourage foundation-model providers from engaging with the healthcare sector or participating in voluntary programs FDA is considering.

V. Any Program for Sharing Foundation-Model Information Must Be Voluntary, Appropriately Scoped, and Protective of Confidential Information (Question 25)

FDA should evaluate generative AI-enabled medical devices principally on the safety and effectiveness of the finished device as configured, validated, and deployed for its intended use. The relevant question is whether the device performs safely and effectively in its intended clinical context, including through the application-layer controls, workflow design, human oversight, validation methods, and postmarket monitoring established by the device manufacturer. Information about the safety measures or technical characteristics of an underlying foundation model should not be presumed to establish, or substitute for, evidence of the finished device's safety and effectiveness. Foundation-model "safety" and medical-device safety and effectiveness are distinct concepts. An upstream model's architecture, training approach, guardrails, or other technical features may not directly predict the performance or clinical risk of a particular device that incorporates or interacts with that model.

FDA should ensure that its review of a finished medical device does not become a vehicle for imposing inappropriate requirements on general-purpose foundation models that are not themselves medical devices. In particular, FDA should not condition device review on comprehensive disclosure of underlying-model information; before requiring particular information about such models, the Agency



should identify why such particular information is necessary to evaluate a specific safety, effectiveness, cybersecurity, or material-change question concerning the finished device.

The Chamber has concerns about whether a Foundation Model Device Master File framework is well suited to this use case. A Master File paradigm may suggest that information about an upstream foundation model can be assessed once and then relied upon across multiple downstream medical-device applications, despite the fact that clinical performance and risk depend substantially on the particular application, intended use, configuration, safeguards, user workflow, and deployment context. If FDA proceeds with a voluntary mechanism for sharing foundation-model information, participation must remain genuinely voluntary, information requests must be appropriately scoped to the finished device under review, and the absence of a Master File or comparable submission must not be treated as a deficiency or otherwise prejudice the review of a device submission.

If FDA establishes a voluntary mechanism for sharing foundation-model information, CDRH should commit explicitly that a device manufacturer's decision not to reference such information, or a foundation-model developer's decision not to submit it, will not be treated as a deficiency; will not result in requests for information that FDA has not identified as necessary to assess the finished device; and will not, by itself, affect the timeline or outcome of premarket review. FDA should evaluate each device submission based on evidence relevant to the finished device's intended use, performance, and risk controls.

To the extent FDA determines that limited foundation-model information is necessary to resolve a specific question concerning the finished device, that information may be among the most competitively sensitive information a technology company holds. Any voluntary information-sharing mechanism must include strong protections against disclosure of trade secrets and other confidential business information, including disclosure under the Freedom of Information Act, and must include clear restrictions on FDA's ability to share the information with other agencies or third parties without the filer's consent. FDA may consider narrowly tailored third-party assessments where a device manufacturer determines that limited information regarding an underlying model is necessary to support review of the finished device. Any such assessment should focus on the specific question relevant to the finished device and should not become a mechanism for generalized FDA review or certification of a foundation model. Any such third-party assessment should be voluntary. Where a manufacturer elects to use such an assessment, the resulting information should be limited to the specific information necessary to address the finished-device review question and should be protected as confidential commercial information.

VI. FDA Should Gather More Input Before Establishing a Regulatory Framework for Agentic AI Systems (Question 26)



Some generative AI systems are being designed to autonomously plan and execute multi-step tasks, use external tools, and take sequences of actions with limited human involvement at each step. The Chamber agrees with CDRH's implicit recognition that these agentic systems present regulatory challenges that are qualitatively different from those posed by single-function generative AI devices.

The Chamber does not believe it is premature to think carefully about how agentic AI should be regulated — but it is premature to build a formal framework before the technology and its clinical applications are better understood. Rules written now are likely to be a poor fit for where the technology ends up, and poorly calibrated requirements for agentic systems could foreclose beneficial applications before they have a chance to develop.

The Chamber recommends that CDRH hold a dedicated public workshop on agentic AI in medical devices before issuing any guidance on this topic. That workshop should address how the two-axis risk framework applies to multi-step systems, what human oversight requirements are necessary and sufficient at different risk levels, and how to distinguish agentic functions that meet the definition of a medical device from those that do not. The Chamber also encourages CDRH to coordinate with other relevant federal agencies, including ONC and CMS, to ensure that any framework for agentic AI in healthcare is coherent across agencies and does not create conflicting obligations for manufacturers, providers, and other stakeholders.

In parallel with a dedicated workshop, the Chamber encourages FDA to consider leveraging its existing authorities to support structured, agency-supervised testing of novel AI tools in real-world settings, generating the evidence needed to inform sound regulatory policy rather than writing rules in the absence of data. The Chamber recommends that HHS establish a voluntary National Clinical AI Sandbox in which states opt in via agreements with HHS, allowing organizations to deploy autonomous clinical AI under phased clinician oversight, defined safeguards subject to FDA's ongoing oversight, and required real-world evidence reporting. This federal-state model would offer a consistent alternative to the emerging patchwork of state AI laws and generate the evidence needed to inform a permanent regulatory framework.

Participation in any Sandbox should be voluntary; should include appropriate safeguards for patients, clinicians, and participating organizations; should not establish new obligations for entities that do not elect to participate; and should be developed through a transparent, multi-stakeholder process. FDA, HHS, states, clinicians, patient representatives, developers, healthcare systems, and other relevant stakeholders should have the opportunity to inform the program's design.

VII. Regulatory Certainty and Accessibility Are Essential to a Functioning Framework



Chamber members across the medical device, health technology, and investment communities consistently identify regulatory uncertainty as one of the primary barriers to investment in generative AI-enabled medical devices. Companies cannot make the multi-year investments required to develop, validate, and commercialize these products without reasonable confidence that the regulatory pathway is predictable and that the rules will not change materially between the time they begin development and the time they seek market authorization.

The Chamber urges CDRH to provide a clear timeline for issuing authoritative guidance following this comment period, and to provide interim clarity, through Q&A documents, pre-submission meetings, and other mechanisms, while more formal guidance is being developed. Pre-submission meetings should be readily available and responsive, particularly for companies developing novel applications where there is no established precedent to draw on. FDA should coordinate this work across HHS to ensure consistent policies, including with the Centers for Medicare & Medicaid Services, Office of the National Coordinator for Health Information Technology, and the Office for Civil Rights, while clarifying the boundaries applicable to non-device clinical software and avoiding duplicative provider obligations.

Many of the most innovative generative AI-enabled medical device developers are small companies and startups that lack the regulatory affairs infrastructure of large device manufacturers. A framework that is technically sound but practically inaccessible to smaller companies will concentrate the market among incumbents and reduce the diversity of innovation that benefits patients. The Chamber urges CDRH to provide plain-language guidance and worked examples that smaller companies can use without incurring undue costs, and to consider a dedicated support program for companies navigating this framework for the first time.

VIII. Conclusion

The Chamber urges FDA to develop a framework that matches oversight to actual risk, preserves established regulatory pathways, draws on standards and evaluation methods developed collaboratively with industry and recognized standards bodies, and gives manufacturers the clarity they need to invest and innovate. The framework should provide a workable path for lower-risk devices through a performance- and monitoring-forward approach; establish clear and fair rules for managing device updates; focus principally on the safety and effectiveness of finished medical devices while strongly protecting confidential foundation-model information; and allow additional evidence and experience to inform future policy for agentic AI systems. The framework should expressly preserve statutory exclusions, recognize established health system clinical governance, and support manufacturers in providing the information healthcare providers need to validate, monitor, and safely deploy AI-enabled tools.



Getting this framework right matters—for patients who stand to benefit from these technologies, for the companies investing in them, and for the United States’ ability to lead in the responsible development of medical AI. The Chamber appreciates the opportunity to submit these comments and looks forward to continued engagement with CDRH.

Sincerely,

A handwritten signature in black ink, appearing to read "Lexi Branson".

Lexi Branson
Vice President, Health Policy
U.S. Chamber of Commerce