



1919 S. Eads St.
Arlington, VA 22202
703-907-7600
CTA.tech

October 19, 2026

Kyle A. Diamantas, Acting Commissioner
U.S. Food and Drug Administration
10903 New Hampshire Avenue
Silver Spring, MD 20993

Docket Number: FDA-2026-N-7874

Dear Acting Commissioner Diamantas:

The Consumer Technology Association (CTA[®]) appreciates the opportunity to comment on the *Considerations for the Regulation of Generative AI-Enabled Medical Devices: Discussion Paper and Request for Feedback*.

CTA appreciates FDA's decision to engage stakeholders early on these important questions and its stated goal of balancing a nimble, risk-based approach consistent with least burdensome principles against the need to ensure timely access to safe, effective, and innovative medical devices.

Artificial intelligence (AI) or super intelligence (SI) has enormous potential to improve health outcomes, expand access to care, and lower costs. FDA's challenge is to protect patients without creating regulatory barriers that delay beneficial technologies or put innovative companies at a disadvantage.

CTA previously provided FDA with recommendations on generative AI in healthcare in a November 11, 2025 letter. In that letter, CTA urged FDA to prioritize high-risk applications, resolve regulatory ambiguity, standardize model validation and evaluation, offer guidance for postmarket monitoring, leverage existing frameworks, and streamline transparency metrics.

We are encouraged that the Discussion Paper reflects many of these themes, in particular its function-based, risk-proportionate orientation, its commitment to least burdensome principles, its recognition that evaluation methods built for bounded-input software do not translate to GenAI, and its willingness to consider accepting greater premarket uncertainty paired with postmarket monitoring.

Our comments focus on practical approaches that can improve safety and effectiveness while giving developers, healthcare institutions, clinicians, and patients clear expectations.

As North America's largest technology trade association, CTA is the tech sector. Our members are the world's leading innovators – from startups to global brands helping support more than 18 million American jobs. CTA owns and produces CES[®] – the most powerful tech event in the world. CTA is the trade association representing more than 1200 companies in the U.S. technology industry. Eighty percent of CTA companies are small businesses and startups; others are among the world's best-known brands. We provide members with policy advocacy, market research, technical education and standards development.

Consumer Technology Association[®]
Producer of CES[®]

CTA is accredited by the American National Standards Institute as a standards developer. To date, we have published almost 40 digital health standards including six specifically dedicated to health AI. Our latest standard, *Operations and Monitoring of Predictive Health AI Solutions (CTA-2136)*, provides a risk-based framework to help organizations monitor performance, detect model drift, and require greater transparency around performance and model updates. This standard builds on *Performance Verification and Validation for Predictive Health AI Solutions (CTA-2135)*, announced in 2025, which provides guidance for evaluating predictive health AI before deployment.

CTA's Health Division advances consumer-based, technology-enabled health solutions to improve health outcomes and reduce overall health care costs. The Division includes telehealth providers, personal health wearable companies, digital health technology companies, healthcare payers, health systems, and biopharmaceutical innovators. Our members use technology to improve nutrition, fitness, mental health, lifestyle management, care access, care coordination, and more – and they are poised to lead the next wave of American innovation with cutting-edge health technology.

Summary

CTA generally supports FDA's effort to adapt its total product lifecycle approach to generative AI (GenAI)-enabled medical devices. GenAI technology can support more personalized, accessible, and efficient care, but open-ended inputs and outputs, variable responses, third-party foundation models, and frequent changes make some traditional evaluation and monitoring approaches difficult to apply without adaptation.

The central challenge is not simply how to test a GenAI-enabled device before clearance or authorization. It is how to define the function and intended use being regulated, evaluate that function in an appropriate way, and maintain reliable evidence and accountability after deployment across different users, workflows, populations, and model versions. CTA recommends that FDA prioritize a practical, scalable postmarket monitoring framework, supported by industry standards, shared infrastructure, and clear responsibility across the ecosystem.

CTA's principal recommendations are:

- Assess risk based on each clinical function as deployed, considering intended use, users, workflow, autonomy, and potential harm.
- Make the two-axis framework more predictable through illustrative use cases and clearer links between risk, oversight, and evidence expectations.
- Use a proportionate premarket approach, combining fit-for-purpose benchmarking and clinical confirmation where warranted, and build on existing standards and the PCCP framework.
- Prioritize practical postmarket monitoring with common metrics, risk-based re-benchmarking, clear reporting and escalation pathways, and privacy-protective information sharing.
- Clarify lifecycle responsibilities and how material changes to prompts, configurations, and third-party models should be assessed, without inadvertently treating healthcare institutions as device manufacturers.
- Address agentic AI through separate, further stakeholder engagement, while clarifying that agentic systems performing non-device functions remain outside device oversight.

Risk Management

Two-Axis Framework

CTA views the proposed framework as a useful starting point. Its continuous dimensions are preferable to rigid categories that may not reflect the range of GenAI applications. The framework, however, does not yet give stakeholders enough detail to determine where a device would fall in the matrix or what regulatory oversight would apply. We recommend FDA:

- Map the current application landscape. Publish, in partnership with outside experts, an illustrative catalog of representative GenAI use cases (e.g., clinical documentation support, patient education, symptom triage, medication guidance, image interpretation) placed on the framework, with the corresponding expected oversight – outside device regulation, enforcement discretion, general controls, or premarket review – and the general evidence tier. This would make regulatory boundaries more granular and expectations more concrete.
- Show how placement on the matrix translates to obligations. The framework is most useful if each region links to default premarket evidence expectations and postmarket monitoring tiers.
- Commit to updating the framework. Because the use landscape is changing rapidly, FDA should revisit the examples periodically and allow risk placement to move, in either direction, as real-world evidence accumulates, including a practical path to down-classify functions that prove lower risk in practice.

Factors That Should Inform Risk

In addition to the two proposed axes, FDA should explain how it will consider the following factors, individually and in combination:

- Clinical function and intended use, including whether the function measures, analyzes, informs, recommends, prioritizes, diagnoses, initiates, or controls an intervention.
 - Degree of directiveness and autonomy, including whether outputs merely provide information, influence a decision, recommend a specific action, or initiate an action.
 - Severity, likelihood, and reversibility of potential harm, including harm from delay, unnecessary escalation, inappropriate intervention, or overreliance.
 - User and context, including whether the intended user is a patient, consumer, generalist clinician, specialist, or trained operator, and the workflow in which the function is used.
 - Opportunity for independent review, including whether a user can reasonably evaluate the basis for an output and whether meaningful human review occurs before consequential action.
 - Availability and reliability of downstream safeguards, such as confirmation steps, escalation pathways, audit logs, source traceability, or fail-safe behavior.
 - Input and output characteristics, including open-endedness, multi-turn interactions, personalization, multimodality, and susceptibility to prompt injection or misleading inputs.
 - Performance uncertainty and changeability, including dependence on third-party foundation models, retrieval systems, prompts, tools, and deployment configurations.
- Population and setting, including the expected distribution of users and patients, clinically relevant subgroups, and resource-constrained or time-sensitive environments.

FDA should avoid turning every factor into a separate mandatory test or documentation requirement. Instead, the factors should help explain why a particular level of evidence, oversight, or monitoring is appropriate for the specific function.

The Deployed Clinical Function Should be the Unit of Analysis

We generally agree that the clinical function as deployed should drive regulation, because the deployed function determines risk, workflow, outputs, and ultimately the infrastructure needed for postmarket monitoring. GenAI rarely refers to one application of one model; it more often describes various applications of various models for various purposes, frequently within a single product. As such, FDA should:

- Confirm that risk is assessed for each device function as deployed, consistent with the Multiple Function Device Products guidance, rather than for the product as a whole or for the underlying model.
- Explain how products with multiple indications or functions at different risk levels will be treated, including how lower-risk or non-device functions are kept outside the scope of review for higher-risk functions.
- Recognize that evaluating a function across the full clinical pathway – rather than isolated outputs – may require statistical and evaluation approaches that differ from those FDA has traditionally relied on, and be open to such approaches.

Consumer Health Tools

Context, specificity, personalization, presentation, and the likely effect on user behavior all matter. FDA should provide examples that distinguish general educational information, contextualized information, personalized recommendations, and explicit instructions. A disclaimer should not by itself determine risk, but neither should the presence of conversational or personalized language automatically make a function a regulated device.

CTA urges FDA to also clarify how the framework interacts with existing policies for general wellness products, software functions excluded from the device definition, clinical decision support, and patient communication tools. This is especially important for consumer wearables and connected health products that may combine wellness insights with a separately defined medical function. The use of GenAI alone should not trigger FDA regulation. Regulatory oversight should depend on what a product does, the risks it presents, and how it is used, not simply the technology that powers it.

Patient-facing functions may create different risks than clinician-facing functions because users have different training and access to safeguards. However, FDA should not assume that patients are inherently unable to use clinical information safely. Evaluation should consider health literacy, usability, communication, escalation pathways, and the actual intended use. Similarly, the distinction between generalist and specialist clinicians should be considered in light of the function and the extent to which independent clinical judgment is expected.

Premarket Evaluation

CTA supports continued exploration of the competency-based approach described in the paper, including device benchmarking followed by clinical confirmation where appropriate. Evaluating the final user-facing device in a representative configuration is important because performance depends on more than the foundation model alone. The approach should, however, remain risk-based and flexible, with the methods, sample sizes, endpoints, and level of clinical confirmation tailored to the intended use.

Benchmarking can help characterize safety behavior, clinical task performance, communication, robustness, and subgroup performance. To be meaningful, benchmarks should be fit for purpose, prespecified where practicable, representative of the intended use, and linked to performance in real-

world conditions. Public benchmarks can be affected by contamination, saturation, and limited representativeness. Sponsor-developed benchmarks can address product-specific risks, provided the methods and acceptance criteria are justified and the limitations are disclosed. Independent evaluation can add value in appropriate cases, but should not be mandatory for every function.

Clinical confirmation should be selected according to the risk and evidence needs of the specific device. Retrospective evaluation, shadow deployment, standardized patient interactions, independent clinician review, and prospective studies may each be appropriate in different circumstances. FDA should make clear that a prospective clinical study is not the default where a less burdensome method can adequately answer the safety and effectiveness questions. Conversely, where the potential harm is serious or existing evidence cannot resolve important uncertainty, more rigorous clinical evidence may be warranted.

For open-ended outputs, FDA should support fit-for-purpose comparators. Depending on the function, a comparator might be a validated reference standard, expert consensus, usual care, unaided clinician performance, or performance of a human-AI team. The comparator should reflect how the device is intended to be used and the clinical question that matters. A single universal comparator, such as specialist performance, may not be appropriate across all use cases.

Synthetic data may supplement real-world data where access is limited or where rare scenarios and underrepresented subgroups need to be explored. It should not automatically be treated as interchangeable with real patient data. Sponsors should explain how synthetic data were generated, what real-world distributions they represent, and how differences and limitations were assessed. Synthetic data should not be the sole basis for concluding that a device performs adequately in a subgroup when the synthetic generation process may reproduce the same blind spots as the system under evaluation.

CTA encourages FDA to work with accredited standards development organizations and other qualified stakeholders to develop reusable evaluation methods, terminology, and reporting practices. Common methods can improve consistency and reduce repeated, product-specific reinvention. Standards should remain adaptable as the technology evolves, and FDA should avoid mandating a particular benchmark or third-party provider where equivalent evidence can support the same conclusion.

Postmarket Monitoring

CTA members view postmarket monitoring as the most important component of a genAI regulatory framework – it is how developers demonstrate safety and effectiveness in real-world use – and also the most challenging. Neither the Discussion Paper nor current industry proposals articulate a robust, practical postmarket system. The agency should consider creating a public-private advisory group to determine how to manage postmarket expectations of evolving technologies. We offer the following recommendations.

Establish Common Monitoring Objectives, Metrics, and Triggers

FDA should work with industry, clinicians, healthcare institutions, patients, standards bodies, and other relevant agencies to define a core set of monitoring concepts and metrics that can be tailored to specific intended uses. These may include:

- Safety-critical errors and events, including failure to recognize escalation needs, materially incorrect recommendations, and failures of intended safeguards.

- Performance degradation or drift associated with changes in users, populations, workflows, data sources, model components, or deployment settings.
- Reliability of refusal, deferral, uncertainty communication, and scope maintenance in clinically relevant scenarios.
- Performance across relevant subgroups and settings, where the data can support meaningful analysis.
- Material changes to the device, including changes to prompts, retrieval sources, guardrails, orchestration, user interfaces, third-party models, or connected tools.
- Use-related signals, including recurring user workarounds, automation bias, overrides, near misses, and patterns of inappropriate reliance.

FDA should distinguish between routine monitoring, a signal requiring investigation, a validated safety concern, and an event requiring urgent action or formal reporting. Clear thresholds and escalation pathways would help prevent both under-reporting and unstructured over-reporting. The framework should not require every output variation or low-severity anomaly to be reported as a safety event.

Comments on Specific Monitoring Approaches

- Periodic re-benchmarking. CTA supports this approach, with cadence scaled to risk and reassessment triggered by defined events: changes to the underlying model, changes to prompts or guardrails that affect safety behaviors, detection of a signal, or a significant shift in the input population. Automation should be encouraged.
- Sample-based clinician review. This approach may be effective, but manufacturers often lack access to real-world interaction data after deployment because of HIPAA, business associate agreements, customer contracts, and data-minimization commitments. FDA should recognize these constraints, allow privacy-preserving and deployer-assisted designs, and avoid expectations that require manufacturers to collect patient data they otherwise would not hold.
- Performance degradation monitoring. CTA supports drift monitoring using defined metrics and proxies, recognizing that all relevant information is often unavailable in real time.
- Additional approaches. FDA should also recognize in-product user feedback and flagging, telemetry on safety behaviors (such as escalation and refusal rates), complaint handling, and registries maintained by professional societies.

Clarify Responsibility Across the Lifecycle

Responsibility should be allocated according to each actor's role, access to information, control over the relevant system, and ability to mitigate the risk. Manufacturers should remain accountable for the safety and effectiveness of their devices, the accuracy of their labeling and intended use, appropriate change control, and the monitoring activities within their control. Deploying institutions should be responsible for deployment decisions and local practices within their control, such as configuration, user training, workflow integration, access management, and escalation processes. Clinicians and users should have clear instructions about intended use and limitations, without being made responsible for identifying technical failures they cannot reasonably detect.

FDA should clarify how these responsibilities apply when a healthcare institution customizes prompts, templates, retrieval sources, or workflows. Not every local configuration should be treated as a change in intended use or as making the institution a device manufacturer. Guidance should distinguish routine configuration within validated boundaries from changes that materially alter clinical functionality or risk. It should also explain how accountability works when a third-party foundation model provider makes a change that the device manufacturer did not initiate.

A clear responsibility map should be established at both the regulatory and operational levels. It should not diffuse manufacturer accountability, but it should recognize that manufacturers may not have direct access to all real-world data and that deployers may be best positioned to observe some use-related signals. Practical guidance should specify what information should be shared, through what channels, and within what timeframes.

Predetermined Change Control Plans

CTA supports building on the existing PCCP framework before creating a separate GenAI-specific regime. The current approach provides a useful foundation, but FDA should clarify how PCCPs can accommodate changes that cannot be fully prespecified in the same way as more bounded software modifications. The goal should be to define the types of changes, methods, acceptance criteria, and safeguards that can be managed within an approved plan, while reserving additional review for changes that materially affect safety, effectiveness, or intended use.

FDA should distinguish among: (1) routine maintenance and low-risk changes; (2) changes within validated parameters and an authorized PCCP; (3) changes that require additional verification or clinical confirmation; and (4) changes that materially alter the intended use or risk profile and may require a new submission. The categories should be illustrative and supported by practical examples, not treated as rigid labels separate from context.

Agentic AI

CTA appreciates the agency including discussion of agentic AI as an emerging issue but we believe any further discussion on this topic should be separate. Agentic systems vary substantially in their level of autonomy, tools, functions, and consequences of action. FDA should not rush to establish new requirements for agentic AI before the technology and its applications are better understood. We recommend a separate discussion paper and dedicated stakeholder process to evaluate the risks, benefits, and appropriate scope of oversight. In the interim, FDA should confirm that agentic systems performing non-device functions, such as documentation, scheduling, coordination, and outreach, remain outside device oversight.

We also recommend that the Department of Health and Human Services (HHS) consider coordinating with states to establish an interim policy framework in which states opt in through agreements with HHS, to gather real-world evidence on agentic AI by allowing the responsible deployment of autonomous clinical AI tools under phased clinician oversight, conditional and revocable FDA enforcement discretion, and required performance 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 that discussion paper and a permanent regulatory framework.

Conclusion

GenAI holds tremendous potential to improve the quality, accessibility, and affordability of health care. It is helping empower users with practical data and insights and is enabling patient advocacy, which can result in better health outcomes, including through discussions with clinicians. Realizing that potential requires a framework that is risk-based, predictable, built on consensus standards, and grounded in a practical, shared approach to postmarket monitoring.

CTA appreciates FDA's willingness to engage stakeholders before establishing its regulatory approach. The discussion paper correctly recognizes that these technologies require thoughtful adaptation of existing regulatory approaches. The most immediate opportunity is to create a clear, scalable framework for postmarket monitoring.

We encourage FDA to continue working with industry, clinicians, healthcare institutions, patients, standards development organizations, and other public-sector partners. A coordinated approach can strengthen safety and trust while preserving the ability to improve products, address emerging risks, and bring beneficial and transformative technologies to patients and clinicians in a timely manner.

FDA has an opportunity to establish a regulatory approach that protects patients, encourages innovation, and strengthens American leadership in health AI. Getting this right means focusing oversight on real risks, building on existing standards, and giving developers clear, practical rules. CTA looks forward to working with FDA to make that happen.

Sincerely,

René Quashie
Vice President, Digital Health
Consumer Technology Association

Kerri Haresign
Senior Director, Technology & Standards
Consumer Technology Association

Catherine Pugh
Director, Digital Health
Consumer Technology Association