



February 23, 2026

Via electronic submission at <http://www.regulations.gov>

Thomas Keane, MD, MBA
Assistant Secretary for Technology Policy,
National Coordinator for Health Information Technology
Department of Health and Human Services
Hubert H. Humphrey Building
200 Independence Avenue, SW
Washington, DC 20201

Re: Request for Information: Accelerating the Adoption and Use of Artificial Intelligence as Part of Clinical Care; 90 Fed. Reg. 60,108 (December 23, 2025)

Dear Dr. Keane,

As the national representative of more than 1,000 leading tax-paying hospitals and health systems throughout the United States, the Federation of American Hospitals (FAH) appreciates the opportunity to submit comments to the Assistant Secretary for Technology Policy (ASTP)/Office of the National Coordinator for Health Information Technology (ONC) regarding the above-referenced *Request for Information: Accelerating the Adoption and Use of Artificial Intelligence as Part of Clinical Care*. We commend the Department of Health and Human Services' (HHS) focus on accelerating the adoption and use of artificial intelligence (AI) specifically within clinical care. FAH members are at the forefront of responsibly integrating AI to enhance clinical care, support the health care workforce and improve system-wide efficiency. When deployed thoughtfully, AI can equip health care providers with better information, reduce administrative burdens, and help direct limited resources where they are needed most. Federal policy should acknowledge and reinforce this leadership by establishing clear, practical guardrails that encourage innovation while safeguarding patients and maintaining public trust. Our comments below offer recommendations to achieve these objectives.

In considering the development and use of AI tools in health care, it is important to first understand how AI technology is defined, and particularly whether a tool incorporates generative AI and machine learning (ML) approaches. Generative AI and ML systems adapt by learning from the data they process, whereas rule-based tools operate predictably and are commonly used in health care to support evidence-based clinical decision-making or clearly defined operational tasks. We note that the Food and Drug Administration (FDA) has authorized over 1,300 devices with AI aspects for marketing in the United States, with AI serving either as a stand-alone product or support other medical technologies (such as AI embedded in a drug pump or surgical robot).

Advances in generative AI and ML are creating opportunities and challenges. **FAH and its members appreciate the promise of AI and the need to carefully balance oversight to ensure safe and appropriate development and use with the need to innovate and continue to advance this transformational technology.** If developed and used responsibly, AI has the potential to revolutionize health care delivery, including improving patient care, operational efficiency, and health outcomes, along with countless other advancements in the field.

We emphasize the unique attributes of health care (e.g., size of the industry; impact on health and safety; health data privacy) and the needs of the multiple health sector actors that will interact with AI technologies, such as AI developers, health care technologists, health care organizations, health care providers and patients/consumers. We also note the need to consider the topic of liability, which is a new and challenging aspect of AI. While health care providers bear

responsibility for the care they provide, the developers of commercial AI products must also be accountable if safety, bias, or other harms are caused by the AI tool itself or a flaw in the tool's development.

We appreciate the RFI's emphasis on identifying actions that HHS can take to "accelerate the adoption and use of AI as part of clinical care," and with that tenet in mind, we make the following key recommendations, among others, to accomplish that goal, as further discussed below:

- **Establish a Uniform Regulatory Framework:** FAH strongly recommends a singular federal regulatory framework, including a national federal privacy law, that expressly preempts state laws, for the use of AI in health care. A patchwork of individual state AI regulations severely limits the ability to deploy AI at scale in health care and limits the potential positive impact AI could have on a health care system's performance.
- **Create a Coordinated, Risk-Management Federal Approach for Clinical AI:** FAH supports a uniform, risk-based approach to AI implementation and oversight, allowing hospitals the flexibility to assess internal risk and implement protections accord to scale, and with a focus on incorporating the existing risk management and security framework already established in health care under HIPAA and HITECH.
- **Modernize Payment Policy to Support AI-Enabled Clinical Care:** The absence of stable, payment pathways is a major barrier to adoption for many AI tools, particularly when embedded within clinical infrastructure rather than offered as discrete, separately billable services, and payment for AI software is insufficient for broader adoption. Thus, payment policy is needed to encourage competition and innovation, especially for AI tools meeting evidence and performance standards.
- **Provider Collaboration:** FAH encourages HHS to institutionalize mechanisms for ongoing provider engagement in clinical AI policymaking,
- **Other Important Areas for Consideration:** liability for AI tools; recognition of permissible clinical AI uses; patient transparency and consent balanced with AI workflow; and research to support workforce AI readiness.

Establish a Uniform Regulatory Framework

Regulatory fragmentation remains the most significant barrier to responsible AI adoption in clinical care and the United States has the opportunity at the federal level to establish a uniform and practical framework to promote AI in health care, considering the context of existing laws and regulations that can be leveraged as a baseline to ensure safe and effective use of AI in health care. **Where new laws and regulations are necessary, we strongly recommend a singular federal regulatory framework, including a national federal privacy law, that expressly preempts state laws, for the use of AI in health care. We caution against a patchwork of individual state AI regulations, as this severely limits the ability to deploy AI at scale in health care and limits the potential positive impact AI could have on an overall health care system's performance.** As acknowledged by the December 11, 2025, [Executive Order](#) directing development of a minimally burdensome national policy framework for AI, the current framework of individual state AI regulations is extremely unwieldy and actively hampers the ability of health systems to leverage AI in an efficient and/or large scale manner. A single federal regulatory framework would go a very long way in alleviating these difficulties. A more uniform approach to AI regulation would catalyze rather than hamper AI development and adoption in health care.

We also urge establishment of standard definitions and benchmarking tools to support effective implementation of AI for health care. For example, the National Institute of Standards and Technology (NIST) has a long history of working across the public and private sectors to develop consensus-based, workable standards and tools to advance the use of technologies that improve the lives of individuals. **We recommend that NIST develop a single set of definitions within the area of AI that can be leveraged by others, avoiding contradictory legal and regulatory approaches. In addition, NIST has the appropriate technical knowledge and experience to develop benchmarking metrics and other tools that both AI developers and deployers would be able to leverage as they seek to ensure that AI solutions are safe, secure, trustworthy, and fit to purpose in health care.** For example, NIST could lead development of standards for model testing, validation, benchmarking and lifecycle monitoring of AI

technologies used in health care. We recommend maintaining a database of validated data sets used for AI technologies to facilitate greater explainability for these technologies.

Create a Coordinated, Risk-Management Federal Approach for Clinical AI

FAH supports a uniform, risk-based approach to AI implementation and oversight. Risk management is a key aspect of ensuring that AI solutions, generative and rules-based, are appropriately developed, disseminated, and monitored over time. For AI solutions in particular, a risk management approach can help both developers and health care providers efficiently focus technical and organizational controls on higher risk deployment.

Risk management approaches are deeply integrated with health care systems, including both existing workflows and regulatory schemes. Hospitals and health care systems have extensive experience in, and have long deployed, risk management approaches to ensure the safety of health care services and the privacy and security of health information.

At the federal level, the existing risk management landscape includes a range of safety and privacy requirements, such as the Medicare Conditions of Participation and Health Insurance Portability and Accountability Act (HIPAA). In addition, health care technologies have established risk management for electronic medical record (EMR) and health record (EHR) workflows. Any AI regulatory requirements that conflict with existing risk management processes will slow down progress in realizing the benefits of technology and could inadvertently result in less effective risk management of complex health care systems and organizations. **An AI regulatory framework should take into consideration the extent to which the AI model directly impacts patient care and should focus on processes to ensure algorithms are transparent, auditable, ethical, fair and safe by incorporating the existing risk management framework already established in health care.** To that end, as discussed above, it is critical that hospitals and health care systems are involved in the creation of AI policies for health care, including establishing standards around risk categories for AI use cases.

Further, while the Food and Drug Administration (FDA), Centers for Medicare & Medicaid Services (CMS), ASTP/ONC and Office for Civil Rights (OCR) each play an essential role in protecting patient safety, supporting interoperability, ensuring appropriate payment, and safeguarding privacy, the absence of coordinated risk-based guidance creates uncertainty for health systems attempting to deploy AI at scale.

Health systems routinely face unresolved questions regarding when AI functionality embedded in clinical workflows constitutes a regulated medical device, how ongoing model updates or retraining affect regulatory status, and where accountability resides when AI tools are developed by vendors but configured, implemented, and monitored locally. This uncertainty slows adoption and can result in overly cautious implementation decisions that limit patient benefit. For this HHS-led effort, we encourage the development of coordinated, cross-agency guidance that clearly explains how existing regulatory frameworks apply to common clinical AI use cases, such as documentation assistance, predictive risk models, clinical prioritization, and diagnostic support. Importantly, such guidance should distinguish between tools that are of high risk of impacting patient safety and those that are of low risk to patient safety and calibrate oversight accordingly.

Leverage Existing HIPAA Frameworks for Clinical AI Data Governance

FAH strongly supports leveraging existing health care privacy and security frameworks to govern clinical AI. The HIPAA Privacy and Security Rules and subsequent Health Information Technology for Economic and Clinical Health Act (HITECH) already provide a comprehensive foundation for managing patient data, and duplicative or parallel AI-specific privacy regimes risk creating confusion and unnecessary burden without improving patient protection. If HHS decides to pursue prescriptive mandates in future regulations regarding transparency, disclosures, and opt-out rights for developers and deployers, we recommend excluding HIPAA Covered Entities from such requirements by adding a robust and comprehensive safe harbor provision. Transparency expectations should be calibrated to clinical risk and operational feasibility, avoiding requirements that would necessitate dual workflows or conflict with established clinical processes.

Clarify Federal Guidance for Non-Medical Device AI in Clinical Care

Many AI applications used in clinical care fall outside traditional medical device regulation, creating novel legal and implementation challenges for providers and vendors. **HHS should not broadly regulate all non-medical device AI**

but should play a coordinating and enabling role by issuing clear, non-binding guidance that explains how existing federal frameworks apply across common deployment models. Such guidance should clarify accountability expectations without prescribing rigid responsibility splits, including considerations related to documentation, ongoing monitoring, model updates and change management, incident reporting boundaries, and security obligations. Consistent with a “One HHS” approach, HHS should align expectations across OCR, ONC, CMS, FDA, and relevant research agencies, promote standardized best practices for procurement, monitoring, and equity assessment, fund shared evaluation and governance infrastructure, and publish voluntary but influential model contracting and accountability guidance to support safe, scalable, and equitable use of non-medical device AI in clinical care.

Explicitly Recognize the Full Scope of Permissible Clinical AI Uses

As AI adoption expands in health care, federal policy should clearly reflect the full range of ways AI is used in clinical and operational settings. Many current and emerging applications extend beyond traditional predictive analytics and require explicit recognition to avoid unintended regulatory gaps. Generative AI models increasingly are used to support clinical documentation by drafting clinical impressions, reports, and summaries that are reviewed, edited, and finalized by clinicians. **Regulatory frameworks should explicitly recognize the creation and drafting of clinical content as a permissible and appropriate use of AI when coupled with clinician oversight, rather than limiting policy constructs to pattern recognition or prediction alone.**

Clinical AI applications also rely on multi-modal inputs, including audio and video. Audio inputs support ambient documentation and workflow efficiency, while video inputs enable safety-focused use cases such as virtual sitters and fall risk monitoring. References to “data sets” in regulatory language may be ambiguous as applied to live or near-real-time feeds. HHS should establish clear guidance clarifying that appropriately governed processing of audio and video inputs is permissible to support patient safety and modern care delivery. Finally, policy language should avoid constraining AI use to narrow concepts such as “assisting treatment” alone. In practice, AI supports a broad operational scope that directly enables clinical care, including triage, patient identification, background safety monitoring, and early risk detection. Regulatory frameworks that recognize this broader context will better support scalable, responsible AI adoption without forcing artificial workflow separations that reduce efficiency or compromise care quality.

Collaborate with Providers to Enable Safe and Scalable AI Adoption

FAH member experience consistently demonstrates that effective AI policy cannot be developed without deep engagement from health care providers who are responsible for delivering care, managing data regulated under HIPAA and other health data privacy and security laws, and integrating technology into complex clinical workflows. As AI moves from pilot use into routine clinical operations, provider engagement must extend beyond consultation toward sustained operational partnership. Health systems bring critical expertise regarding workflow design, clinician adoption, patient communication, and unintended consequences that may not be evident in premarket evaluation alone. **We encourage HHS to institutionalize mechanisms for ongoing provider engagement in clinical AI policymaking,** including participation in pilots, demonstrations, and post-deployment monitoring initiatives, such as FDA post-market surveillance. These partnerships will help ensure that federal guidance is feasible across diverse care settings and aligned with the realities of clinical practice.

Balance Patient Transparency with Integrated Clinical AI Workflows

As HHS considers approaches to patient transparency, engagement, and trust in clinical AI, it is essential to account for the integrated nature of many AI applications within health care delivery. Proposals that rely on individualized opt-out mechanisms for AI systems, while well-intentioned, may unintentionally undermine both care quality and operational feasibility. Many AI tools used in clinical care are deeply embedded within electronic health record systems and care delivery platforms, enhancing safety and efficiency in real time. For example, AI-enabled clinical decision support may flag potential drug–drug interactions or surface risk indicators as clinicians are actively ordering medications or managing care. In these contexts, enabling individual opt-outs would require the creation of parallel non-AI workflows to replicate the same functions through manual review. Such parallel systems are technically complex, costly to maintain, and may increase the risk of error rather than reduce it.

Clinical experience also demonstrates that AI can improve diagnostic accuracy and timeliness, particularly in time-sensitive and resource-constrained settings. AI-assisted radiology tools, for example, can rapidly identify imaging patterns consistent with acute stroke, supporting faster diagnosis and treatment in emergency departments where specialist availability may be limited. Requiring human-only review or bifurcated workflows based on individual

preference could delay care and negatively affect outcomes, even for patients who might otherwise choose to opt in but are unable to do so because of their clinical condition.

The technical and practical challenges associated with individualized opt-out requirements, including software redesign, dual workflow maintenance, and additional staff training, risk deterring health care organizations from adopting AI solutions altogether. In some cases, limiting AI use could also result in worse patient outcomes where no viable human substitute exists. Maternal-fetal monitoring provides one illustrative example: AI-enabled systems could continuously analyze high-frequency physiologic data to identify emerging risk patterns and support earlier clinical intervention, capabilities that are difficult to replicate reliably through manual review alone.

For these reasons, **we encourage HHS to avoid policies that rely on individualized opt-out mechanisms for integrated clinical AI systems.** A more balanced approach would emphasize transparency about how AI is used in care delivery, incorporate general consent for AI use within care settings consistent with existing workflows for efficiency purposes, and rely on robust oversight mechanisms to ensure safety, efficacy and equity through continuous monitoring of model performance and outcomes.

Modernize Payment Policy to Support AI-Enabled Clinical Care

In the Hospital Outpatient Prospective Payment System for Calendar Year 2026 Final Rule, CMS acknowledged that rapidly emerging software technologies, including software-as-a-service (SaaS) tools designed to support clinical decision-making, do not yet fit cleanly within existing payment frameworks. The absence of stable, repeatable payment pathways remains a major barrier to adoption for many AI tools—particularly those embedded within clinical infrastructure rather than offered as discrete, separately billable services.

Under the Inpatient Prospective Payment System, the [New Technology Add-on Payment \(NTAP\)](#) has provided a limited pathway for reimbursement of certain AI software, but its current scope is narrow and insufficient to support broader deployment. Similarly, while the American Medical Association has established [Current Procedural Terminology \(CPT\) codes](#) for some AI applications, these codes largely focus on diagnostic testing and do not address many documentation, workflow, and decision-support use cases. Clear guidance on coverage expectations and medical record documentation requirements for these technologies would support more consistent adoption. Additionally, moving from fragmented [Category III CPT codes](#) toward more stable, valued reimbursement pathways for AI-augmented analysis would further support responsible adoption.

To encourage competition and innovation, payment policies should be structured so that any AI tool meeting the same evidence and performance standards may qualify under the same payment framework, regardless of whether it is a vendor or development model. This approach would avoid preferential treatment and support both commercial and internally developed tools that meet clinical and operational standards. In addition, HHS should continue to play a central role in governing the use of AI in medical necessity determinations and utilization management, ensuring that such tools protect patient access, reduce administrative friction, and align with established coverage principles. Lastly, AI-related costs should be recognized as allowable expenses within value-based payment models, coupled with incentives for outcomes-based contracting. Payment policies should avoid penalizing providers for earlier or more accurate identification of high-risk conditions enabled by AI.

Clarify Developer and Deployer Expectations Without Redefining Liability

As AI adoption accelerates, it is essential that federal policy clearly distinguish between AI developers and AI deployers. Health care providers frequently deploy third-party AI tools within clinical environments and may also develop internal tools that are not commercialized. Federal guidance should recognize this reality and avoid categorizing health care providers as developers subject to obligations intended for entities that design, commercialize, and market AI products.

At the same time, deployment of AI in clinical care should reinforce the principle that AI is intended to support—not replace—clinical judgment. Consistent with emerging federal guidance, including themes articulated in the [Centers for Medicare & Medicaid Services AI Playbook](#), AI tools should be implemented with appropriate human-in-the-loop governance. Clinicians must remain the final decision-makers and validators of AI-informed outputs, whether those outputs relate to care recommendations, documentation support, clinical test interpretation, diagnostic assistance, or workflow prompts.

Maintaining human oversight reflects longstanding clinical accountability models in which clinicians are responsible for patient care decisions regardless of the tools used to support those decisions. AI tools should augment human work by improving efficiency, accuracy, and situational awareness while preserving clinician accountability. Clear role definitions and human-in-the-loop expectations also align with payment and compliance considerations. As federal payment systems evolve to accommodate software-enabled technologies, billed services must continue to reflect properly delivered and documented care, regardless of whether AI tools are used to support clinical workflows. Reinforcing these principles supports patient safety, program integrity, and responsible innovation without imposing developer-level obligations on health care providers.

Shared Responsibility and Developer Accountability

There is a shared responsibility between the developers and deployers/end-users of AI tools to build and deploy them in a way that is safe, effective, and secure. **While health care providers bear responsibility for the care they provide, the developers of commercial AI products must be accountable for the safety and reliability of their products and required to be truthful in marketing their products, especially since safety, bias, privacy and security, or other harms may be caused by a flaw in the tool itself.** Commercial AI and machine learning developers must provide deployers/end-users of their tools with guidance on ethical use, such as when it is necessary to have “a human in the loop,” and the limits of their models. Deployer/end-users will also need guidance on how to provide oversight of AI tools that are in use to ensure that they are functioning appropriately over time, along with collaboration from developers in properly monitoring AI technology performance over time.

Adopt an Evidence-Centric Lifecycle Approach to Clinical AI Evaluation

HHS should adopt an evidence-centric, lifecycle approach to evaluation of clinical AI that governs how tools are validated, monitored, and maintained over time. This approach should include rigorous multisite pre-deployment validation, standardized and transparent reporting artifacts, and instrumented continuous monitoring post-deployment to assess performance, safety, and equity as models evolve. Public clinical AI test beds and pragmatic trials should be used to support real-world validation across diverse care settings and patient populations, complemented by independent test laboratories capable of generating reproducible and comparable evidence. Standardized methods for robustness testing, drift monitoring, bias and fairness assessment, human factors and usability, and workflow impact should be applied consistently across the lifecycle.

To operationalize this framework, HHS should leverage policy and programmatic levers—including grants, Centers for Medicare & Medicaid Services Innovation Center pilots, Cooperative Research and Development Agreements (CRADAs), and support independent evaluation infrastructure—to embed these requirements into practice. In doing so, HHS should build on and recognize mature, provider-tested evaluation, accreditation, and governance approaches already developed and in use across the private sector, rather than introducing entirely new frameworks that would require broad retooling by providers and vendors. Aligning lifecycle evidence generation with payment policy will be essential to rewarding clinical value, discouraging premature deployment, and enabling scalable, responsible adoption. In addition, HHS should enable and recognize robust, voluntary private accreditation and certification frameworks for clinical AI that are standards-based and updated as technologies evolve. Recognition of standardized test beds, credentialing programs, and evaluation processes—supported by public reference datasets, aligned federal guidance, and appropriate procurement incentives—would help create trusted market signals without imposing prescriptive regulatory requirements.

Prioritize Implementation-Focused Research and Workforce Readiness

To accelerate responsible adoption of AI in clinical care, HHS should prioritize research investments that strengthen both evaluation infrastructure and workforce-centered implementation science. HHS should fund regulatory science and develop a shared evaluation infrastructure to support the responsible adoption of clinical AI, including the creation of clinical AI test beds for real-world validation and investment in high-quality, ethically sourced, representative AI-ready datasets. Use of CRADAs to co-develop reusable platform components and common evaluation tools would help establish durable measurement standards that promote transparency, unlock competition and reduce contracting and implementation costs.

Research involving nurses and the nursing workforce is a critical area that the HHS should prioritize to accelerate the adoption of AI in clinical care. Nurses represent one of the largest segments of the United States health care workforce and are foundational to care delivery across virtually all clinical settings. Numbering more than five million professionals across registered nurses, licensed practical and vocational nurses, and advanced practice registered nurses, the

nursing workforce plays a central role in health care delivery. While AI research in other clinical domains offers substantial promise, research focused on understanding adoption drivers, workforce impacts and implementation challenges specific to nursing is essential for scalable and sustainable deployment.

Nurses also serve as primary adopters and sustainers of clinical technologies. Over time, they have integrated increasingly complex devices and digital tools into care environments while maintaining the human-centered aspects of care. As AI-enabled technologies continue to expand, research designed by and for nurses will be critical to ensuring these tools enhance care quality and efficiency without introducing unintended burden. In alignment with HHS's goal of designing infrastructure responsive to user needs, priority research areas should include nurses' perspectives on AI, baseline understanding of AI-enabled tools, drivers and barriers to adoption across care settings, and strategies to support safe, compliant integration of AI into evidence-based nursing workflows. Research examining both the benefits and risks of AI in specific nursing contexts will be essential to supporting responsible adoption and strengthening public trust.

FAH appreciates the opportunity to address these critical issues, and looks forward to collaborating with ASTP/ONC in advancing the use of AI for the benefit of patients, providers, communities and federal health care programs. As leaders in thoughtfully deploying AI, our member tax-paying hospitals can play a pivotal role with HHS as it develops and implements AI regulatory clarity, patient-centered transparency, payment policy and implementation-focused research while leveraging existing privacy and compliance frameworks. This RFI represents a meaningful opportunity to move from high-level AI principles to actionable, clinically grounded policy. If you have any questions or would like to discuss further, please do not hesitate to contact Katie Tenover, SVP & General Counsel, at ktenover@fah.org.

Sincerely,
/s/
Charlene MacDonald
President and CEO