HEALTH INNOVATION ALLIANCE

440 1st Street NW, Suite 430, Washington, DC 20001 | 202.643.8645 | Health-Innovation.org

July 9, 2026
International Medical Device Regulators Forum
Artificial Intelligence/Machine Learning-enabled Working Group
c/o IMDRF Secretariat
Health Sciences Authority of Singapore (2026 Chair and Secretariat)
Submitted electronically via the IMDRF AIML Working Group Secretariat

RE: Public Consultation on IMDRF/AIML WG/N93 DRAFT: 2026 – Technical Framework for Artificial Intelligence Life Cycle Management

To the Artificial Intelligence/Machine Learning-enabled Working Group:

The Health Innovation Alliance (HIA) is pleased to submit comments on the International Medical Device Regulators Forum (IMDRF) draft technical framework, Artificial Intelligence Life Cycle Management (IMDRF/AIML WG/N93 DRAFT: 2026). We appreciate the Working Group's continued leadership on harmonizing international considerations for AI-enabled medical devices, building on the Good Machine Learning Practice (GMLP) guiding principles (IMDRF/AIML WG/N88 FINAL: 2025) and the broader body of IMDRF work on software as a medical device.

HIA is a diverse coalition of patient advocates, healthcare providers, consumer organizations, employers, technology companies, and payers working together to improve health care through the commonsense use of data and technology. Founded in 2007, HIA has been a consistent voice for patient-centric, pro-innovation health policy, including on federal AI frameworks in the United States.

Summary of HIA's Position

We are generally supportive of the draft framework's central premise: that AI-enabled medical devices benefit from a harmonized life cycle approach grounded in quality management, risk management, human oversight, and cybersecurity. In particular, we agree with the document's recognition that core principles of medical device life cycle management (systematic planning, requirements management, traceability, risk management, and validation) remain essential for AI-enabled medical devices rather than a parallel regulatory structure.

HIA's comments focus on areas where the draft's current framing risks undermining the risk-based, use-case based approach that the document otherwise endorses. AI in health care varies enormously across use-cases, patient populations, clinical environments, and degrees of autonomy. Oversight expectations that are calibrated to this variability will foster innovation while protecting patients; expectations that apply uniformly regardless of risk will impose compliance burdens disproportionate to benefit and may push manufacturers away from jurisdictions that adopt the framework most aggressively.

Our detailed comments, set out in the accompanying table, address nine specific concerns. Three themes run through them:

- **Consistency on risk-based framing.** Several sections of the draft (notably the references to "continuous" post-market monitoring, the universal framing of human oversight, and the detailed logging list in Section 5.6) are in tension with the risk-based approach articulated in Section 4.2. We recommend aligning these sections with the document's own risk framing.
- **Device-level validation as the unit of analysis.** Validation of AI-enabled medical devices should focus on the device in its intended use environment. Model-level external validation is a useful input, but it is not the appropriate regulatory endpoint. Transparency and validation expectations for third-party general-purpose models and foundation models should reflect what manufacturers reasonably control.
- **Architecture and supplier neutrality.** The framework should be neutral between deployment architectures (including cloud and local hosting) and should clearly allocate compliance responsibility to the legal manufacturer, consistent with existing IMDRF practice. Expectations on third-party suppliers, cloud providers, and foundation model providers should be bound by the scope of control each entity has over the device.

A risk-based, context-dependent approach, applied consistently across the life cycle, will allow the framework to achieve its harmonization goals without dampening innovation or disadvantaging manufacturers that build on modern infrastructure.

Detailed Comments

Nine specific comments, cross-referenced to the relevant sections and line numbers of the draft, are provided in the table on the following page. We would welcome the opportunity to discuss any of these comments with the Working Group and to serve as a resource as IMDRF finalizes the technical framework.

Thank you for the opportunity to comment on this important harmonization effort.

Sincerely,

Roslyn Docktor
Executive Director
Health Innovation Alliance

cc:
Michelle Tarver, M.D., Ph.D., Director, Center for Devices and Radiological Health, U.S. Food and Drug Administration
Richard Abramson, Director, Digital Health Center of Excellence, U.S. Food and Drug Administration

ANNEX: DETAILED COMMENTS

IMDRF/AIML WG/N93 DRAFT: 2026

Technical Framework for Artificial Intelligence Life Cycle Management

Submitter: Health Innovation Alliance (HIA), Washington, DC, United States

Contact: Roslyn Docktor, Executive Director, rdocktor@health-innovation.org

No.	Section	Line(s)	Comment and Rationale	Proposed Change
1	4.4 Cybersecurity	343-353	The draft implies that cloud connectivity introduces greater cybersecurity risk and offers local hosting as a mitigation ("limiting the frequency of updates to batches, hosting the AI-enabled medical device locally, if possible"). This framing is not architecture-neutral. Modern cloud architectures can provide equal or greater security than local hosting through centralized patching, continuous monitoring, hardware security modules, and defense-in-depth controls that are difficult to replicate on-premises. Cybersecurity expectations should be met through appropriate design and controls regardless of deployment model. IMDRF guidance that suggests a preferred architecture risks steering manufacturers away from deployment choices that may be more secure and more scalable in practice, particularly for smaller manufacturers and lower-resourced health systems.	Revise Section 4.4 to remove the suggestion that local hosting is a preferred mitigation. Replace with architecture-neutral language: cybersecurity controls (encryption, access control, monitoring, patch management, anomaly detection) should be applied proportionate to the risks identified regardless of whether the device is hosted locally, on a manufacturer server, or in a cloud environment. Reference IMDRF/CYBER WG/N60 FINAL:2020 as the governing framework rather than prescribing a deployment model.

No.	Section	Line(s)	Comment and Rationale	Proposed Change
2	5.2 Data Collection and Management	464-519; 690-692	Section 5.2 treats detailed training data transparency as a baseline expectation for all AI-enabled medical devices, and Section 5.3 expects "transparent information about the use of general-purpose and off-the-shelf models." This framing does not adequately account for the growing use of foundation models and third-party general-purpose models in which detailed training data information is proprietary, unavailable to the device manufacturer, or commercially sensitive. Requiring or strongly suggesting disclosure of underlying training data for these models would disadvantage manufacturers that build on third-party infrastructure and would not meaningfully improve safety, since the appropriate assessment point for an AI-enabled medical device is its validated performance in its intended context of use, not the upstream training corpus. Transparency is a legitimate goal, but expectations should be calibrated to what the manufacturer reasonably controls and to what actually informs safe and effective use.	Revise Section 5.2 and the general-purpose model subsection of Section 5.3 to clarify that: (a) transparency expectations apply to information the manufacturer reasonably controls or can obtain from its suppliers; (b) where training data details for third-party or foundation models are proprietary or unavailable, the manufacturer should document this limitation, characterize the implications for validation, and demonstrate fit-for-purpose performance through device-level validation rather than upstream data disclosure; and (c) proprietary model information should not be required to be disclosed publicly, consistent with international practice on protecting trade secrets while supporting regulatory assessment.

No.	Section	Line(s)	Comment and Rationale	Proposed Change
3	5.4.1 V&V Activities related to the Model	746-756	The document identifies "External Validation of the Model" as a discrete life cycle activity distinct from clinical evaluation of the AI-enabled medical device. While separating model-level and device-level activities can aid clarity, the current framing risks entrenching model-in-isolation validation as an expectation, which is not the appropriate unit of regulatory analysis. The model is a component; the medical device in its intended use environment is the regulated product. Conflating validation of the model with validation of the device could lead to duplicative evidence generation, regulatory divergence across jurisdictions, and a misallocation of manufacturer resources away from the clinical performance questions that actually determine patient benefit and risk.	Revise Section 5.4.1 to clarify that external validation of the model is one input to device-level verification and validation, not a standalone regulatory endpoint. The overarching unit of analysis for AI-enabled medical devices should be the device in its intended use environment, consistent with the SaMD clinical evaluation framework in IMDRF/SaMD WG/N41 FINAL:2017 and the device-level validation approach taken by FDA. Section 5.4.1 should be framed as supporting the clinical evaluation described in Section 5.4.2, not as a parallel requirement.
4	5.6 Operations and Monitoring	970-975	Section 5.6 recommends that logging mechanisms capture detailed logs and audit trails of "model predictions and any explainability elements (if available) such as heatmaps or top predictive features, data inputs, model outputs, and metadata such as timestamps, user interactions, model versions, and environmental context." While logging is important for higher-risk devices and adaptive models, this level of detail is not appropriate as a baseline expectation for all AI-enabled medical devices regardless of risk. Comprehensive logging of explainability elements, user interactions, and environmental context for lower-risk FDA designated Class I and Class II devices would generate substantial data infrastructure, storage, and privacy burdens disproportionate to the risk profile, and would create friction for the small and mid-size manufacturers IMDRF aims to support.	Revise Section 5.6 to frame the list of logging elements as considerations that may be appropriate for higher-risk, adaptive, or generative AI-enabled medical devices, with lower-risk devices expected to maintain logging proportionate to their classification and clinical context. Align the section with the risk-based framing in Section 4.2 and with existing IMDRF risk characterization work (IMDRF/SaMD WG/N81).

No.	Section	Line(s)	Comment and Rationale	Proposed Change
5	Cross-cutting (Sections 4.4, 5.1, 5.6, 5.7)	346; 452-454; 983-986	Multiple sections of the document reference "continuous" post-market monitoring and surveillance as an expectation for all AI-enabled medical devices (for example, "continuous post-market monitoring and surveillance" in Section 4.4 and Section 5.6, and "continuous monitoring of the ongoing performance" in Section 5.1). This is inconsistent with the risk-based approach that the document otherwise endorses. Not all AI-enabled medical devices, particularly those cleared through FDA 510(k)-equivalent pathways, warrant continuous post-market monitoring plans. The inconsistency between the document's risk-based framing (Sections 4.2, 5.4) and its blanket references to continuous monitoring could lead to fragmented implementation across jurisdictions and unnecessary burden on lower-risk devices.	Replace blanket references to "continuous" monitoring with risk-proportionate language throughout the document. A consistent formulation, such as "ongoing monitoring commensurate with the device's risk classification, adaptivity, and context of use," would align these sections with the risk-based framework established in Section 4.2 and with the variability in post-market expectations across existing jurisdictional frameworks.
6	4.3 Human Oversight	291-310	Section 4.3 describes human oversight as "essential throughout the entire life cycle" of AI-enabled medical devices. While human oversight is critically important in many clinical contexts and particularly in decision support, framing it as universal and life-cycle-wide conflicts with the reality that certain autonomous and semi-autonomous AI-enabled medical devices are specifically designed to operate with reduced or episodic human involvement, and are authorized on that basis in some jurisdictions. Treating human oversight as a universal design requirement could constrain innovation in autonomous diagnostics, closed-loop therapeutic systems, and other devices where reduced human-in-the-loop involvement is a clinical and regulatory feature rather than a deficiency.	Revise Section 4.3 to frame human oversight as context-dependent and risk-proportionate. A preferable formulation: "Human oversight is an important consideration throughout the AI-enabled medical device life cycle. The nature, frequency, and life cycle points at which human oversight is required should be determined based on the device's intended use, risk profile, and degree of autonomy, and should be consistent with the regulatory authorization in the relevant jurisdiction."

No.	Section	Line(s)	Comment and Rationale	Proposed Change
7	4.4 Cybersecurity; Introduction to Life Cycle	354-365	Section 4.4 uses "primary responsibility" language when describing the legal manufacturer's obligations for validation of input data and references "additional parties such as end users and sites" as also being involved. Although well-intentioned, this framing introduces ambiguity about whether third parties (including cloud infrastructure providers, data suppliers, and deploying sites) carry secondary regulatory obligations. In most jurisdictions, regulatory compliance responsibility rests with the legal manufacturer of the AI-enabled medical device. Blurring this allocation complicates contracting, liability, and audit expectations, and creates friction for manufacturers that rely on third-party infrastructure (including cloud services and foundation models) to build safe and effective devices.	Revise Section 4.4 to clarify that compliance responsibility for the AI-enabled medical device rests with the legal manufacturer. Third-party suppliers (including cloud providers, data suppliers, foundation model providers, and deploying sites) may be engaged through appropriate supplier management practices and contractual arrangements, and audit expectations for these entities should be limited to the scope of control they have over the device's performance. This aligns with existing IMDRF supplier management guidance referenced in IMDRF/SaMD WG/N23 FINAL:2015.
8	5.3 Model Building and Tuning (Generative AI); 5.6 Operations and Monitoring	667-693; 1019-1029	The subsections addressing general-purpose, off-the-shelf, and generative AI-enabled models are underdeveloped relative to the pace of technology adoption. The risk profile for devices incorporating large language models and other foundation models varies widely by use case, and a number of the recommendations, such as collecting "interaction logs and prompt-response histories from LLM's" and monitoring "prompt chaining and workarounds" for all generative AI-enabled devices, are not appropriate as uniform expectations. Lower-risk administrative and documentation use cases warrant lighter monitoring than clinical decision support or autonomous generative systems. A more granular, risk-based treatment would better serve both manufacturers and regulators.	Expand the generative AI and general-purpose model subsections to articulate a tiered, risk-based framework that distinguishes administrative, documentation, decision-support, and autonomous use cases. Align monitoring and logging expectations for generative AI-enabled medical devices with the risk-proportionate framing recommended in Comments 4 and 5 above. Reference GMLP Guiding Principles and emerging international standards (including ISO/IEC 23053:2022 and ISO/IEC 42001 on AI management systems, see Comment 9) to support this framing.

No.	Section	Line(s)	Comment and Rationale	Proposed Change
9	3 References; 4.1 Quality Management System	141-192; 199-219	The document references ISO 13485 and a range of AI-specific standards (ISO/IEC 23053, ISO/IEC 25058, ISO/IEC 25059, ISO/IEC 5259-4) but does not reference ISO/IEC 42001:2023, the international standard on AI management systems. ISO/IEC 42001 is increasingly being adopted by technology companies and would provide a coherent complement to ISO 13485 for manufacturers building AI management practices on top of their existing quality management systems. Incorporating a reference to ISO/IEC 42001 in the Quality Management System section (4.1) and in the References list would help align IMDRF's framework with the broader AI governance standards landscape.	Add ISO/IEC 42001:2023 (Information technology – Artificial intelligence – Management system) to Section 3 References. In Section 4.1, note that manufacturers may consider implementing AI management practices consistent with ISO/IEC 42001:2023 alongside their ISO 13485-compliant quality management system.

RELEVANT HIA PUBLICATIONS

Our recent work in this space includes the HIA Principles for the Use of Artificial Intelligence in Health Care and Life Sciences (2024), the Use Cases for AI Tools to Improve Health Care report (2025), HIA's response to the White House Office of Science and Technology Policy Artificial Intelligence Action Plan Request for Information (March 2025), HIA's comments on FDA's draft guidance on the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products (April 2025), and HIA's response to the National Institutes of Health Artificial Intelligence Request for Information (July 2025). They are linked below and are available at https://health-innovation.org:

1. Health Innovation Alliance, Principles for the Use of Artificial Intelligence in Health Care and Life Sciences (July 2024).
2. Health Innovation Alliance, Use Cases for AI Tools to Improve Health Care (March 2025).
3. Health Innovation Alliance, Response to the Office of Science and Technology Policy Request for Information on the Development of an Artificial Intelligence Action Plan (March 13, 2025).
4. Health Innovation Alliance, Comments to the U.S. Food and Drug Administration on Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products (April 7, 2025).
5. Health Innovation Alliance, Response to the National Institutes of Health Request for Information on Artificial Intelligence (July 15, 2025).