

Electronic Health Record Developer Adherence to Federal AI Risk Transparency Requirements

Nidhi Manikoth

Georgetown University

Ngan Vo

Georgetown University

Jessica Handley

MedStar Health National Center for Human Factors in Healthcare

Josh Biro

MedStar Health National Center for Human Factors in Healthcare

Jordan Everson

Georgetown University

Raj M. Ratwani

Raj.M.Ratwani@medstar.net

MedStar Health National Center for Human Factors in Healthcare

Kristen Miller

MedStar Health National Center for Human Factors in Healthcare

Short Report

Keywords:

Posted Date: December 12th, 2025

DOI: <https://doi.org/10.21203/rs.3.rs-8255773/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Version of Record: A version of this preprint was published at npj Digital Medicine on October 14th, 2026. See the published version at <https://doi.org/10.1038/s41746-026-03152-w>.

Abstract

We evaluated electronic health record developer compliance with federal transparency requirements for AI-enabled predictive decision support interventions. Developers are required to publicly report on risk analysis, mitigation, governance, and on eight other attributes. Across 35 developers, reporting was inconsistent and often incomplete, with substantial variation in structure and detail, yet all developers were certified as meeting the requirements. Results suggest a need for standardized reporting and more rigorous certification review.

Introduction

Artificial intelligence (AI) algorithms are increasingly embedded in electronic health records (EHRs) to support diagnostic reasoning, risk prediction, clinical decision-making, workflow automation, and treatment recommendations.¹ These tools hold promise for enhancing patient safety by improving diagnostic consistency, reducing preventable errors, and providing timely clinical insights.² Beyond clinical decision support, emerging high-value uses of AI within EHRs include reimbursement optimization, operational management, and quality and safety monitoring.³ However, concerns have emerged regarding model bias, data quality, explainability, inappropriate generalization, and their potential to introduce new forms of patient safety risks.⁴ Studies have documented that AI-enabled decision support can propagate or amplify inequities when underlying training data are unrepresentative, lack important clinical context, or reflect structural biases in care delivery.^{5–7} Additional risks—including inadequate validation, lack of transparency in model provenance, variable performance across demographic groups, and unclear lines of accountability—pose challenges for clinicians, patients, regulators, and developers.^{8–10}

To begin to address some of these concerns, the U.S. Department of Health and Human Services Assistant Secretary for Technology Policy (ASTP)/Office of the National Coordinator for Health Information Technology (ONC) sought to bring greater transparency around EHR developer practices related to AI algorithms embedded in their products.¹¹ These regulations recognize that EHR-integrated algorithms, particularly those offering predictive or diagnostic decision support, require clear communication about their development, intended use, evaluation, and known limitations to support safe and trustworthy adoption.¹²

Under the Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing (HTI-1) Final Rule, certified health information technology (CHIT) developers (e.g. EHR developers) must provide publicly accessible Risk Management Summaries (RMS) for AI algorithms, which are broadly considered predictive decision support interventions (DSIs).¹¹ These summaries are required to describe the developer's risk analysis process, mitigation strategies, and governance oversight. They should also address eight core attributes recommended for evaluating DSIs: validity, reliability, robustness, fairness, intelligibility, safety, security, and privacy. The intent of this

requirement is to enable greater transparency for clinicians, health systems, purchasers, and regulators, and to promote accountability for responsible AI development.¹³

Despite these expectations, little is known about how EHR developers interpret and operationalize the HTI-1 RMS requirements or how consistently they report information in practice. Understanding how developers interpret these requirements is essential because health systems increasingly rely on these tools without visibility into their safety or governance processes. To address this gap, we systematically analyzed all available EHR developer RMS to determine alignment with HTI-1 requirements and identify opportunities for policy refinement and improvements in EHR developer reporting practices.

Methods

Data Source and Sample Identification

Using the publicly available HTI-1 certified health IT product list, we identified all products that included an AI-enabled predictive DSI. For each certified product, the corresponding RMS was downloaded for review. All documents were compiled, cataloged, and prepared for systematic review.

Coding Framework and Procedures

A structured coding framework was developed based on the requirements outlined in ONC's HTI-1 Final Rule. The framework assessed whether each RMS (1) described its risk analysis process, (2) identified specific risks, (3) outlined risk mitigation strategies, (4) described governance structures, and (5) addressed any of the eight recommended attributes for evaluating DSIs: validity, reliability, robustness, fairness, intelligibility, safety, security, and privacy. Conceptual synonyms or closely related terminology (e.g., "bias mitigation" for fairness) were accepted as evidence that an attribute was addressed. Two reviewers were trained on the coding framework and independently coded all RMS documents. Coding included both binary indicators of attribute presence and qualitative notes describing the level of detail provided. Discrepancies were reviewed jointly and resolved through consensus.

Analytic Approach

We conducted descriptive quantitative analyses to summarize the frequency and distribution of reported risk management components across developers. Qualitative synthesis of reviewer notes was used to characterize patterns in reporting practices, variation in completeness and specificity, and common gaps relative to HTI-1 expectations. Because all data were publicly available and contained no identifiable information, this study did not require human subjects review.

Results

A total of 35 developers disclosed at least one AI-enabled DSI in their certified products. Among these, 71% (n = 25) described their overall risk analysis process, though the specificity of these descriptions varied. Only 57% (n = 20) explicitly identified the individual risks they monitor as part of this process. In contrast, most developers (91%, n = 32) outlined one or more mitigation strategies intended to address identified risks, and 80% (n = 28) reported having a formal governance structure in place to oversee risk management activities.

When examining alignment with the eight DSI attributes required by ONC/ASTP, 77% (n = 27) referenced all eight attributes at least once within their RMS. However, the depth and clarity of these references differed considerably across developers, with some providing only brief or high-level statements and others offering more detailed explanations of how the attributes were assessed or addressed.

Discussion

These results demonstrate considerable variability in how EHR developers report their risk management practices for AI enabled predictive DSIs. While most developers attempted to address the elements required by the HTI-1 Final Rule, the scope and specificity of their reporting differed substantially. Some developers focused primarily on listing the risks they had identified while others described portions of the broader lifecycle, such as how risks are identified, how they are analyzed, or how mitigation strategies are implemented. In several cases, RMS documents provided only high-level statements without supporting detail. Despite these incomplete or uneven summaries, all products were certified as meeting the reporting requirements.

The RMS documents varied widely in structure, terminology, and formatting. Developers used different organizational approaches and the level of explanation for required elements ranged from brief references to more extensive narrative descriptions. This inconsistency limits transparency and makes it difficult to compare practices across products or evaluate the robustness of individual developers' risk management approaches. Even when information was provided, it was not always usable by a broad audience limiting the practical utility of these reports in supporting responsible adoption of AI-enabled decision support. To improve accessibility, ASTP/ONC should provide a standardized reporting template or structure for RMS submissions. Likewise, certification bodies responsible for verifying adherence to HTI-1 reporting requirements should ensure that developers provide sufficient and appropriately detailed information. Historically, some certification organizations have not consistently conducted rigorous reviews of developer submissions, which may contribute to the variability seen across RMS documents.¹⁴ EHR developers themselves also have an opportunity to strengthen transparency by providing more complete descriptions of their risk management practices.

Our findings echo concerns raised in recent literature about the current lack of uniformity in CDS governance frameworks and risk-management practices despite growing adoption of AI-enabled CDS in health systems. A recent review article calls for standardized reporting and local oversight to ensure

equitable, safe, and transparent deployment; our observed inconsistencies in publicly available RMS documents indicate these recommendations are not yet realized.¹⁵

This study has limitations, including reliance solely on publicly available RMS documents without access to proprietary or internal developer materials. As a result, the analysis reflects only what developers chose to disclose, not the full extent of their practices. It is possible that some variability reflects early-stage adoption of the RMS requirement, but even so, the inconsistencies highlight a need for clearer expectations and oversight. Nonetheless, the intent of the HTI-1 Final Rule is to foster responsible and safe integration of AI into healthcare by increasing visibility into developer risk management processes.¹² Achieving this goal will require greater standardization in reporting and more rigorous evaluation of RMS submissions to support transparency, accountability, and trust in predictive DSIs.

Declarations

Competing Interests

The authors declare no competing interests.

Funding

Funding: Not applicable.

Author Contribution

NM, NV, RR and KM developed the coding framework. NM and NV performed the coding. NM, NV, and KM adjudicated discrepancies. JB, JH, JE, RR and KM conceptualized the study and reviewed the analyses. NM, NV, JB, JH, JE, RR and KM drafted, revised, and approved the final manuscript.

References

1. He J, Baxter SL, Xu J, Xu J, Zhou X, Zhang K. The practical implementation of artificial intelligence technologies in medicine. *Nature medicine*. 2019;25(1):30-36.
2. Bates DW, Levine D, Syrowatka A, et al. The potential of artificial intelligence to improve patient safety: a scoping review. *NPJ digital medicine*. 2021;4(1):54.
3. Sahni NR, Carrus B. Artificial intelligence in US health care delivery. *New England Journal of Medicine*. 2023;389(4):348-358.
4. Chusteki M. Benefits and risks of AI in health care: Narrative review. *Interactive Journal of Medical Research*. 2024;13(1):e53616.

5. Hussain SA, Bresnahan M, Zhuang J. The bias algorithm: how AI in healthcare exacerbates ethnic and racial disparities—a scoping review. *Ethnicity & Health*. 2025;30(2):197-214.
6. Chin MH, Afsar-Manesh N, Bierman AS, et al. Guiding principles to address the impact of algorithm bias on racial and ethnic disparities in health and health care. *JAMA Network Open*. 2023;6(12):e2345050-e2345050.
7. Chen RJ, Wang JJ, Williamson DF, et al. Algorithmic fairness in artificial intelligence for medicine and healthcare. *Nature biomedical engineering*. 2023;7(6):719-742.
8. Reddy S, Rogers W, Makinen V-P, et al. Evaluation framework to guide implementation of AI systems into healthcare settings. *BMJ health & care informatics*. 2021;28(1):e100444.
9. Fehr J, Citro B, Malpani R, Lippert C, Madai VI. A trustworthy AI reality-check: the lack of transparency of artificial intelligence products in healthcare. *Frontiers in Digital Health*. 2024;6:1267290.
10. Daneshjou R, Smith MP, Sun MD, Rotemberg V, Zou J. Lack of transparency and potential bias in artificial intelligence data sets and algorithms: a scoping review. *JAMA dermatology*. 2021;157(11):1362-1369.
11. Services. USDoHaH. Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing. . 2024.
12. Everson J, Smith J, Marchesini K, Tripathi M. A regulation to promote responsible AI in health care. *Health Affairs Forefront*. 2024;
13. AI N. Artificial intelligence risk management framework (AI RMF 1.0). URL: <https://nvlpubs.nist.gov/nistpubs/ai/nist-ai>. 2023:100-1.
14. Ratwani RM, Benda NC, Hettinger AZ, Fairbanks RJ. Electronic health record vendor adherence to usability certification requirements and testing standards. *Jama*. 2015;314(10):1070-1071.
15. Lin AL, Parrish AB, Cary M, et al. Algorithm-Based Clinical Decision Support: Evolving Regulatory Landscape and Best Practices for Local Oversight. *Annual Review of Biomedical Data Science*. 2025;8