SPECIAL REPORT

UNDERSTANDING ONC'S HEALTH AI TRANSPARENCY AND RISK MANAGEMENT REGULATORY FRAMEWORK

April 2024

McDermott Will & Emery | McDermott+ Consulting

McDermott Will & Emery | McDermott+ Consulting

SPECIAL REPORT

TABLE OF CONTENTS

LEARN MORE

For more information, please contact your regular McDermott lawyer, or:

ALYA SULAIMAN
PARTNER
Asulaiman@mwe.com
Tel +1 310 7886017

JAMES CANNATTI
PARTNER
jcannatti@mwe.com
Tel +1 202 7568866

DANIEL GOTTLIEB
PARTNER
dgottlieb@mwe.com
Tel +1 312 9846471

KAREN SEALANDER
COUNSEL
ksealander@mwe.com
Tel +1 202 7568024

NATHAN GRAY
ASSOCIATE
Ngray@mwe.com
Tel +1 202 7568971

RACHEL STAUFFER
VICE PRESIDENT
rstauffer@mcdermottplus.com
Tel +1 202 2041460

KRISTEN O'BRIEN
VICE PRESIDENT
klobrien@mcdermottplus.com

For more information about McDermott Will & Emery visit mwe.com

INTRODUCTION

The Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing (HTI-1) final rule, issued by the US Department of Health and Human Services (HHS) Office of the National Coordinator for Health Information Technology (ONC), went into effect on March 11, 2024, and establishes a new regulatory framework for certain artificial intelligence (AI) and machine learning technologies that support decision-making in health care. Compliance is required beginning December 31, 2024.

This *Special Report* provides an overview of the final rule's algorithmic transparency and risk management requirements and the types of technologies that may be in scope. ONC released an initial proposal for this health AI regulatory paradigm in its HTI-1 proposed rule in April 2023, which we summarized in a previous *Special Report*. Throughout this *Special Report* we note some key differences between the proposed and finalized certification criterion for decision support interventions (DSIs).

We provided an overview of the final rule's updates to certain standards and other certification criteria in the ONC's Health IT Certification Program (certification program) in a separate *On the Subject*, and discussed the final rule's information blocking provisions in another *On the Subject*.

KEY TAKEAWAYS

- Compared to the proposed rule, the final rule significantly scales back the types of algorithms, models and other technologies subject to the new certification requirements.
- The final rule creates requirements for and a description of "evidence-based" DSIs, which ONC describes as "*only those DSIs that are actively presented to users in clinical workflow to enhance, inform, or influence decision-making related to the care a patient receives and that do not meet the definition for Predictive DSI*." Note, however, that the requirements for evidence-based DSIs are not included in this *Special Report*.
- Predictive DSIs (as defined below) supplied by certified health information technology developers (certified health IT developers) as part of a Health IT Module (*i.e.*, certified health IT software applications or electronic health records) are subject to the final rule's new certification requirements.
- Certified health IT developers with Predictive DSIs that are in scope under the new DSI certification criterion must support and maintain certain technical and quality information (called "source attributes") and implement risk management practices for each Predictive DSI. Compliance is required by December 31, 2024.
- The final rule places the knowledge, decision and ongoing compliance obligations associated with Predictive DSI supplied by a certified health IT developer solely within the control of certified health IT developers, even if the Predictive DSI was developed by another entity.

BACKGROUND

Since 2010, the certification program has maintained a certification criterion for clinical decision support (CDS) functionality integrated into the electronic health record (EHR). In the last 10 years, the clinical decision support landscape has evolved considerably. Predictive models are increasingly used and relied upon to inform many health care decision-makers, including clinicians, allied health professionals, payers, patients and researchers. In response to these advancements, the final rule replaces the previous CDS certification criterion for certified health IT with a new DSI certification criterion, creating new technical, transparency and risk management requirements for a range of contemporary and emerging functionalities and software applications that aid decision-making across all areas of health care, including clinical workflows, billing, scheduling, public health disease surveillance and clinical research. ONC has described this new approach as aligned with the federal government's efforts to promote trustworthy AI.

The new DSI certification criterion establishes a definition for Predictive DSIs and requires certified health IT developers that certify to the criterion and supply Predictive DSI as part of their Health IT Modules to (1) support and maintain source attributes, (2) implement intervention risk management practices and (3) make certain summary information about their risk management practices publicly available for each Predictive DSI.

In describing its policy basis for its rule, ONC identified concerns about growing evidence that predictive models introduce or increase the potential for a variety of risks that create unintended or adverse impacts on patients and communities. These risks can impact health care decisions in myriad ways, including through predictive models

that exhibit harmful bias, are broadly inaccurate, have degraded due to model or data drift, are incorrectly or inappropriately used, or widen health disparities.

Note that the final rule's preamble states that certified health IT developers are not required to have certified health IT applications that supply, enable or interface with certified Predictive DSIs. Those that do, however, must satisfy the new criterion's requirements for such technology beginning on December 31, 2024.

WHAT QUALIFIES AS A PREDICTIVE DSI?

The final rule provides the following definition of predictive decision support intervention:

> **Predictive decision support intervention** or **Predictive DSI** means technology that supports decision-making based on algorithms or models that derive relationships from training data and then produces an output that results in prediction, classification, recommendation, evaluation, or analysis.

- The term "**intervention**" in predictive DSI is not intended to mean an intervention as the term is used in clinical care *(i.e.*, prescribing medicine, a medical procedure or other clinical treatment) but rather, an intervention that occurs within a user's workflow, regardless of the type of user or the workflow context.
- Predictive DSIs can take a variety of forms in workflow, including but not limited to alerts, order sets, flowsheets, dashboards, patient lists, documentation forms, relevant data presentations, protocol or pathway support, reference information or guidance, and reminder messages, if those technologies support decision making based on algorithmis or models that dervice relationships from training data, and produce an output that results in prediction, classification, recommendation, evaluation or analysis.

Technically, ONC declined to limit the definition of Predictive DSIs based on risk, context of use or the specific source or developer of the intervention. ONC also emphasized that "who develops" a Predictive DSI is separate and distinct from how ONC interprets what qualifies as a Predictive DSI. However, as explained below, the final rule is clearly focused on a much narrower range of algorithms and AI models than the proposed rule. Specifically, the final rule focuses the scope of the Predictive DSIs subject to the DSI certification criterion to those Predictive DSIs that are supplied by certified health IT developers as part of their Health IT Modules.

FOCUS ON PREDICTIVE DSI SUPPLIED BY CERTIFIED HEALTH IT DEVELOPERS AS PART OF HEALTH IT MODULES

While the definition of Predictive DSI remains broad, the final rule's DSI certification criterion is functionally limited in its applicability because it focuses several key requirements on only those Predictive DSIs that are supplied by certified health IT developers as part of Health IT Modules (*e.g.*, software applications listed on ONC's certified health IT products list). Health IT Module is defined as any service, component or combination thereof that can meet the requirements of at least one certification criterion adopted by the HHS Secretary.

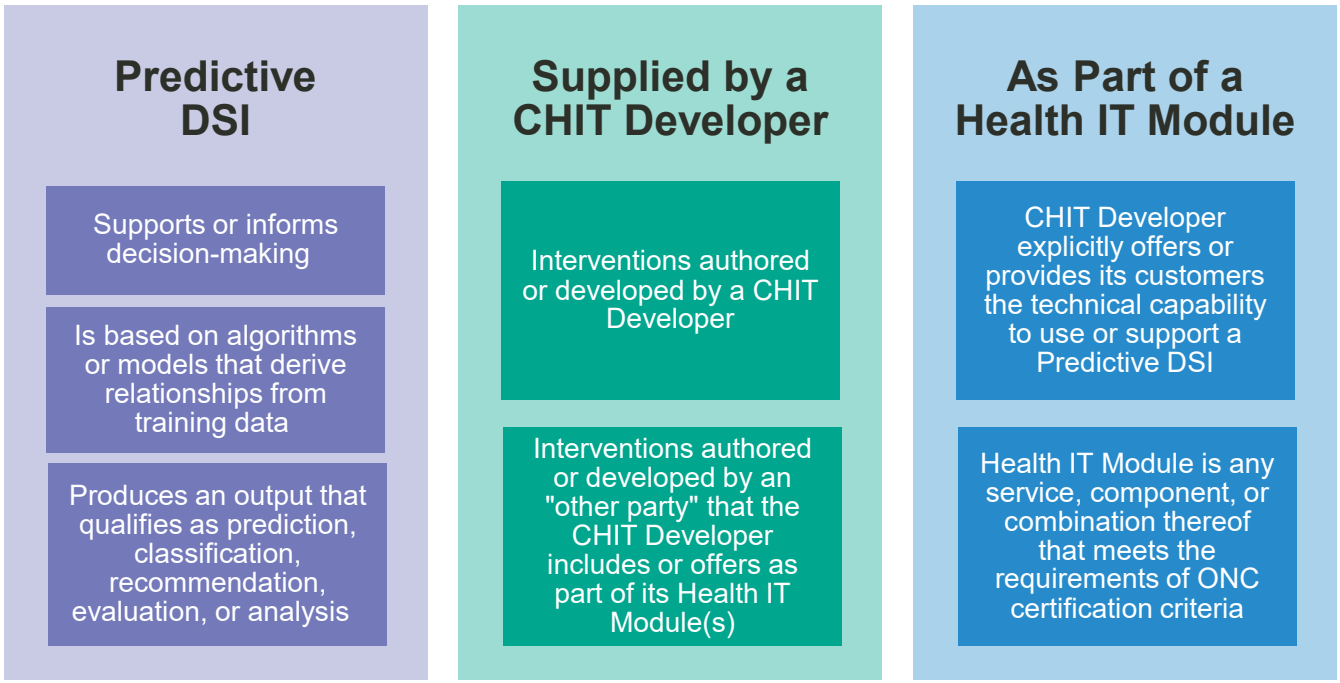

Predictive DSIs "supplied by" a certified health IT developer includes interventions authored or developed by the certified health IT developer as well as interventions authored or developed by an "other party" that the certified health IT developer includes as part of its Health IT Module. For example, if an "other party" enters into a contract with a certified health IT developer to enable EHR end users to access and use a Predictive DSI through the EHR, that "other party-developed" Predictive DSI could fall within the scope of ONC's regulatory framework. ONC states that the concept of "supplied by" is meant to convey that the certified health IT developer has taken on stewardship and accountability for that Predictive DSI for the purposes of its Health IT Module.

A Predictive DSI is supplied "as part of a Health IT Module" when the certified health IT developer has explicitly offered or provided its customers the technical capability to use or support a Predictive DSI, regardless of whether the Predictive DSI was developed by the developer of certified health IT or by an "other party."

The final rule also distinguishes between Predictive DSIs and evidence-based DSIs. Evidence-based DSIs, for purposes of ONC's regulatory requirements, are limited to only those DSIs that are actively presented to users in clinical workflow to enhance, inform or influence decision-making related to the care a patient receives and that do not meet the definition for Predictive DSI. Evidence-based DSIs are subject to a different set of transparency and other requirements compared to Predictive DSI. Requirements for evidence-based DSI are not summarized in this Special Report. Note, however, that the evidence-based DSI definition could implicate many existing clinical decision support software functions within Health IT Modules.

WHAT DOES NOT QUALIFY AS A PREDICTIVE DSI?

Importantly, ONC excluded unsupervised learning models from the scope of Predictive DSI to focus on models trained in data without labels. From ONC's perspective, unsupervised models are generally based on data without

labels, which often generate measures of similarity or closeness of observations rather than a predicted value. In these instances, ONC acknowledges that assessing the accuracy, validity and fairness of a prediction would be difficult, if not impossible, because the outcome is not specified. ONC also clarified that a health care provider who self-develops a tool that meets the definition of Predictive DSI is not subject to the ONC's regulatory requirements.

EXAMPLES OF PREDICTIVE DSI

While the final rule narrows the range of applications subject to the new requirements, the types of technology that meet the Predictive DSI definition is still not obvious. However, in the preamble discussion in the final rule, ONC provides some examples of technologies that would and would not likely be considered Predictive DSI. It is worth noting that many of the technologies ONC describes as likely not Predictive DSI are most similar to unsupervised machine learning models, which ONC explicitly excluded from the scope of the Predictive DSI definition. ONC provides the following examples:

Would *Likely* Be Considered Predictive DSI	Would *Unlikely* Be Considered Predictive DSI
Models that generate clinical notes or draft clinical notes and that were trained based on relationships in large data sets of free text, including large language models (LLMs), and support decision making about what to document in the clinical note.	Indices and classification systems developed by expert consensus rather than in empirical data, such as the SOFA index and NYHA Heart Failure classification, would likely not be considered Predictive DSIs but are likely evidence-based DSI because the score is based on pre-defined rules and not relationships learned in training data.
Models that predict whether a given image contains a malignant tumor or that predict patient reported pain based on an image, trained based on relationships observed in large data sets often using neural networks.	Rules-based algorithms for routing secure messages based on the type of message, rather than relationships in training data.
Models that predict risk of sepsis, readmission (*e.g.,* LACE+), estimated glomerular filtration rate (eGFR) or risk of suicide attempt, which have been trained based on relationships observed in large data sets, often using logistic regression and machine learning techniques, and are used to support decision-making.	Growth charts –for instance percentile calculations based on a lambda-mu-sigma transformation of similar age children's weights – with parameters learned in training data from a national sample of children; because the underlying model is based on the distribution of a single variable *(e.g.*, weight) rather than a prediction based on relationships between variables.
Models that pre-selected or highlighted a default order from an order set based on relationships in training data indicating that such order was most likely to be selected.	A calculation for body mass index (BMI) would likely not be considered a Predictive DSI because the calculation (weight divided by height squared) is not based on relationships in training data.
Models that use natural language processing to route secure messages, which were trained based on the relationship between message contents and the individual who responded to similar messages in the past.	Patient-matching algorithms based on indices of similarities, rather than by relationships in training data where an outcome is known, would likely not be Predictive DSIs.
	Models that use natural language processing to route secure messages, which were trained based on the relationship between message contents and the individual who responded to similar messages in the past.

WHAT DOES ONC REQUIRE FOR HEALTH IT MODULES WITH PREDICTIVE DSI?

Developers and users of Predictive DSIs should take note that HTI-1 places the knowledge, decision and ongoing compliance obligations associated with offering a Predictive DSI solely within the control of certified health IT developers. The final rule imposes several requirements on certified health IT developers providing Health IT Modules with Predictive DSIs. These requirements include:

1. **Enable Users to Activate DSIs:** A limited set of identified users must be enabled to select (*i.e.*, activate) both evidence-based and Predictive DSIs.

2. **Implement Risk Management and Governance Practices:** For each Predictive DSI supplied by the certified health IT developers as part of their Health IT Module, the certified health IT developers must implement risk management practices, including subjecting each Predictive DSI to risk analysis and risk mitigation related to:

   - Validity
   - Reliability
   - Robustness
   - Fairness
   - Intelligibility
   - Safety
   - Security
   - Privacy

   Further, all Predictive DSIs supplied by a certified health IT developer as part of its Health IT Module must have in place policies and controls for governance, including how data are acquired, managed and used.

   Summary information about each of these practices must be made publicly available via a hyperlink before December 31, 2024. This required summary information is far less involved than the detailed documentation requirements in the proposed rule.

3. **Populate and Maintain Source Attribute Information and Functionality:** Health IT Modules must support and maintain a list of "source attributes" (*i.e.*, categories of technical performance and quality information) for both evidence-based and Predictive DSIs. Health IT Modules must also enable several functionalities related to source attributes. The source attributes requirements are discussed in more detail below.

4. **Review and Update Information and Intervention Risk Management Practices:** The source attribution information, risk management practices and provided summary information must be reviewed and updated by certified health IT developers over time.

By December 31, 2024, Health IT developers will need to update their Health IT Modules to meet the DSI criterion's requirements and provide the updated certified health IT to customers. **Starting January 1, 2025**, developers with health IT certified to the DSI criterion must comply with new certification maintenance requirements, and the DSI criterion will become required for health care providers to have certified health IT that

continues to meet the Base EHR definition. Meeting the Base EHR definition is required for providers to be in a position to have "Certified EHR Technology" for purposes of certain Medicare payment programs.

UNPACKING PREDICTIVE DSI SOURCE ATTRIBUTE REQUIREMENTS

Among the most notable features in the final rule's new health AI regulatory framework is the requirement that certified health IT developers supplying Health IT Modules with DSIs maintain and support source attribute data for each evidence-based and Predictive DSI. Source attributes are categories of technical performance and quality information. The list of source attributes (see chart below) in the final rule is lengthy, but ONC emphasized that the required information is not intended to include detailed data on model parameters or hyperparameters or specifics around how input or output variables are defined, transformed or otherwise operationalized.

Certified health IT developers are responsible for supporting source attributes for any Predictive DSI in their Health IT Modules regardless of whether the certified health IT developer developed the Predictive DSI itself or if it was developed by an "other party" that contracted with the certified health IT developer. Beginning January 1, 2025, certified health IT developers must adhere to ongoing maintenance requirements to keep DSI "source attribute" information complete and up to date.

Source Attributes for Predictive DSIs

Category	Details	
I. Details and output of the intervention	1.	Name and contact information for the intervention developer
	2.	Funding source of the technical implementation for the intervention(s) development
	3.	Description of value that the intervention produces as an output
	4.	Whether the intervention output is a prediction, classification, recommendation, evaluation, analysis or other type of output
II. Purpose of the intervention	5.	Intended use of the intervention
	6.	Intended patient population(s) for the intervention's use
	7.	Intended user(s)
	8.	Intended decision-making role for which the intervention was designed to be used/for (*e.g.,* informs, augments, replaces clinical management)
III. Cautioned out-of-scope use of the intervention	9.	Description of tasks, situations, or populations where a user is cautioned against applying the intervention
	10.	Known risks, inappropriate settings, inappropriate uses or known limitations
IV. Intervention development details and input features	11.	Exclusion and inclusion criteria that influenced the training data set
	12.	Use of variables related to race, ethnicity, language, sexual orientation, gender identity, sex, date of birth, social determinants of health data, and health status as input features
	13.	Description of demographic representativeness according to variables related to race, ethnicity, language, sexual orientation, gender identity, sex, date of birth, social

Source Attributes for Predictive DSIs		
Category	**Details**	
		determinants of health data, and health status including, at a minimum, those used as input features in the intervention
	14.	Description of relevance of training data to intended deployed setting
V. Process used to ensure fairness in development of the intervention	15.	Description of the approach the intervention developer has taken to ensure that the intervention's output is fair
	16.	Description of approaches to manage, reduce or eliminate bias
VI. External validation process*	17.	Description of the data source, clinical setting, or environment where an intervention's validity and fairness has been assessed, other than the source of training and testing data
	18.	Party that conducted the external testing
	19.	Description of demographic representativeness of external data according to variables related to race, ethnicity, language, sexual orientation, gender identity, sex, date of birth, social determinants of health data, and health status including, at a minimum, those used as input features in the intervention
	20.	Description of external validation process
VII. Quantitative measures of performance	21.	Validity of intervention in test data derived from the same source as the initial training data
	22.	Fairness of intervention in test data derived from the same source as the initial training data
	23.	Validity of intervention in data external to or from a different source than the initial training data*
	24.	Fairness of intervention in data external to or from a different source than the initial training data*
	25.	References to evaluation of use of the intervention on outcomes, including, bibliographic citations or hyperlinks to evaluations of how well the intervention reduced morbidity, mortality, length of stay, or other outcomes*
VIII. Ongoing maintenance of intervention implementation and use	26.	Description of the process and frequency by which the intervention's validity is monitored over time
	27.	Validity of intervention in local data*
	28.	Description of the process and frequency by which the intervention's fairness is monitored over time
	29.	Fairness of intervention in local data*
IX. Updated and continued validation or fairness assessment schedule*	30.	Description of the process and frequency by which the intervention is updated
	31.	Description of the frequency by which the intervention's performance is corrected when risks related to validity and fairness are identified

**The Health IT Module must indicate when this information is not available for review.*

In addition to supporting the source attributes listed above, certified health IT developers must enable certain functionality related to the source attributes. Health IT Modules must support the technical capability for *other party* source attribute information to be entered into the listedsource attribute fields. However, certified health IT developers are not required to receive, acquire, or otherwise obtain source attribute information for an *"other party's"* Predictive DSI (unless such Predictive DSI is *supplied by* a certified health IT developer as part of its Health IT Modules). A certified health IT developer could include capabilities for other parties to record source attributes in a way that shields the certified health IT developer from accessing the information.

Required Source Attribute Functionalities for CHIT Developers/Health IT Modules with DSIs

At minimum, have the technical capability for users and other parties to populate source attributes themselves	Enable users to identify source attributes and record, change, and access those source attributes based on local validation
Enable a limited set of identified users to record, change, and access additional source attribute information not specified in the rule	Enable users to access emerging transparency measures specific to emerging Predictive DSI types, such as those based on LLMs

Health IT Modules must also enable a limited set of identified users to access complete and up-to-date descriptions of all source attributes related to evidence-based DSIs and Predictive DSIs that are supplied by the certified health IT developer as part of their Health IT Modules. The reference to a "limited set of identified users" means the required capabilities can be constrained to a smaller userbase identified as having the privileges necessary to access and use the capabilities.

Health IT Modules must also enable a limited set of identified users to record and change source attributes, including accessing additional source attributes not specified in the chart above. Nothing in ONC's proposal would require a user of a Health IT Module to review or otherwise engage with source attribute information, although ONC notes that certain users may already have an existing obligation to review source attributes to ensure compliance with non-discrimination requirements and comply with applicable law.

OVERLAP WITH FDA SOFTWARE AS A MEDICAL DEVICE (SAMD) FRAMEWORK

In the final rule, ONC emphasized that its adopted approach is harmonized with and complementary to the framework of the Food and Drug Administration (FDA) for evaluating SaMD applications. ONC stated that where the ONC's and FDA's regulatory oversight responsibilities intercept, ONC supports a consistent approach to regulating predictive technology, independent of the platform on which such technology operates. That said, the FDA and ONC have separate and distinct authorities and regulate for purposes and policy objectives that are distinct from one another. But ONC did highlight similarities between its requirements and those the FDA imposes on certain medical devices. For example, in a scenario where device-Clinical Decision Support Software has been cleared, approved or otherwise authorized for marketing by the FDA, ONC states that the device's manufacturer will have ready access to much of the information necessary for it to comply with ONC's requirements (if the device manufacturer is also a certified health IT developer).

CONCLUSION

The final rule is a major development in the regulation of AI in health care and has implications for organizations throughout the health and information technology industries. Certified health IT developers and other parties that may contract with certified health IT developers to integrate or offer health AI tools into EHRs should familiarize themselves with the definition for Predictive DSI and the transparent and risk management requirements that apply to Predictive DSI.

Please do not hesitate to contact the authors or your regular McDermott lawyer if you have questions or need assistance with understanding the impacts of the HTI-1 final rule.

AUTHORS

ALYA SULAIMAN
PARTNER

asulaiman@mwe.com
Tel +1 310 7886017

JAMES CANNATTI
PARTNER

jcannatti@mwe.com
Tel +1 202 7568866

DANIEL GOTTLIEB
PARTNER

dgottlieb@mwe.com
Tel +1 312 9846471

KAREN SEALANDER
COUNSEL

ksealander@mwe.com
Tel +1 202 7568024

NATHAN GRAY
ASSOCIATE

ngray@mwe.com
Tel +1 202 7568971

KRISTEN O'BRIEN
VICE PRESIDENT

klobrien@mcdermottplus.com
Tel +1 202 756 8964

RACHEL STAUFFER
VICE PRESIDENT

rstauffer@mcdermottplus.com
Tel +1 202 2041460

McDermott Will & Emery | McDermott+ Consulting

mwe.com