

HTI-5 and the New Ground Rules for Health Data: What the Comment Letters Actually Say

MAR 16, 2026





Table of Contents

What HTI-5 Is and Why It Matters Now

The Certification Cleanup: Mostly Agreed, Deeply Cautious

The Information Blocking Wars: Where the Real Fight Is

AI Transparency Gets Gutted and Nobody Is Happy About It

The Automation and RPA Question That Could Break Everything

What This Means for Entrepreneurs and Investors

Abstract

HTI-5 is the fifth iteration of ONC's Health Data, Technology, and Interoperability rulemaking. Published December 29, 2025, comment period closed February 27,

Core proposals: (1) Remove 34 of 60 EHR certification criteria, revise 7 more, totaling nearly 70% of existing requirements. (2) Update information blocking definitions: narrow exception pathways. (3) Lay groundwork for a FHIR-forward, API-first certification ecosystem with room for agentic AI.

Projected savings: \$1.53B total, 1.4M compliance hours in year one, avg 4,000 hrs/developer.

Key commenters in this analysis: Epic, EHR Association (~27 member companies: Oracle Health, Altera Digital Health, MEDITECH, PointClickCare, Premier Inc, NCQA, MGMA (70K+ members, 350K+ physicians), American College of Physicians (163K members), Allina Health, AHIP (205M covered lives), HL7 International, Wolters Kluwer, Innovaccer. More were submitted but not analyzed.

Industry fault lines: Certification cleanup = broad support. Information blocking exception changes = EHR vendors strongly opposed. AI transparency rollback = mostly negative. RPA/agentic AI write access = some support (data intermediaries), strong opposition (EHR platforms, clinical orgs). TEFCA exception removal = broadly supported.

Investor/entrepreneur implications: Significant. Lower barriers to certification enable API-first competitive dynamics, data intermediary tailwinds, AI clinical decision support market reset, information blocking litigation risk redistributed.

What HTI-5 Is and Why It Matters Now

ONC released HTI-5 right at the end of 2025, framed explicitly as a Trump EO 13771 deregulatory play. On paper it is a cleanup operation – decades of accumulated certification requirements that nobody defends in public but everyone has built compliance workflows around. Underneath that cleanup is a fairly aggressive set of bets about where health IT infrastructure is headed, and those bets touch almost every investment thesis in the sector.

The short version: ONC wants to strip out legacy, document-centric, functionality-specific cert criteria and replace the entire program foundation with FHIR-based standards. They also want to close some loopholes in the information blocking framework that, in their telling, have allowed EHR incumbents to use regulatory exceptions as a moat. These are two very different things packed into one proposed rule, and the comment letters treat them almost as separate documents – universal support for the cert cleanup in principle, and a messy, legally charged fight over information blocking pieces.

For anyone building in or investing around EHR connectivity, clinical AI, data interoperability, or health data infrastructure, this rule matters more than most IT policy in the last several years. The certification program is the regulatory scaffolding that determines what features certified EHR products are required to have, which shapes what APIs exist, what data you can access and how, and what competitive dynamics look like between EHR platform owners and third-party vendors. Changing 70% of it in one rule is a big move.

Before getting into what the industry actually said, two contextual points are worth noting. First, the comment period closed February 27, 2026, and a final rule has yet been issued. Many of the most controversial proposals will be modified or dropped before finalization. Second, this rule is simultaneously a deregulatory effort and, for certain market participants, a significant increase in regulatory exposure. That tension runs through almost every substantive comment letter.

The Certification Cleanup: Mostly Agree, Deeply Cautious

The broad industry consensus is that removing outdated certification criteria is a good idea, but that ONC has not thought carefully enough about sequencing, transition timelines, and downstream effects on the providers who actually use this stuff.

The criteria being removed fall into roughly three buckets. The first bucket is genuinely obsolete – stuff like the Clinical Quality Measures Filter criterion, which was designed for a CMS program that was retired before it ever launched, or the Consolidated CDA Creation Performance criterion, which Epic flatly described as having never added value. Nobody fights hard for these and they get removed with minimal drama in the comment letters.

The second bucket is more complicated: functionality-focused document exchange criteria, particularly anything related to C-CDA creation and the Direct Project secure messaging standard. ONC wants to remove requirements to certify the ability to send C-CDA documents, reasoning that industry has moved on to FHIR-based

APIs. The EHR Association, Oracle Health, Altera, MEDITECH, and virtually every clinical organization pushed back hard on this. Oracle's analytics showed its customer base generating roughly 40 million C-CDA documents per month. Carequality, one of the major national exchange networks, reportedly facilitates over 1.2 billion documents exchanged monthly. Direct Secure Messaging has cumulatively processed over 6.5 billion messages since tracking began. This is not a niche legacy capability being quietly phased out – it is the operational backbone of a substantial portion of the current health data exchange. The argument from these commenters is not that C-CDA should exist forever, but that removing the certification requirement before FHIR document standards reach equivalent maturity creates a transition valley where the old thing stops being reliably certified and the new thing is not yet ready. HHS made this point clearly, noting that the FHIR R4 Document specification is still in Trial Use status and that full implementation guides replacing C-CDA for all document types do not yet exist.

The third bucket is where it gets interesting for entrepreneurs and investors: criteria that cover clinical decision support AI transparency, privacy and security controls, safety-enhanced design testing, and real-world testing requirements. These get their own sections below because they are where the most commercially significant fights are happening.

On implementation timelines, almost every commenter flagged the proposed effective dates as too aggressive. The two options ONC proposed are the date of the final rule or January 1, 2027, and neither gives sufficient time for CMS to update its Promoting Interoperability program requirements to stay in sync with what certified EHR technology is actually required to do. Altera, Oracle, and the MGMA all made this point in detail – providers are being held to CMS reporting requirements that depend on certification criteria ONC is proposing to remove, and the coordination between the two agencies appears inadequate. This is not just an administrative inconvenience. If a practice's Promoting Interoperability measures depend on a certified C-CDA workflow and the certification requirement disappears while the CMS requirement stays, that practice is stuck in a compliance gap with real payment consequences.

The Information Blocking Wars: Where the Real Fight Is

This is the section where the regulatory environment gets genuinely contentious where the entrepreneurial implications are largest. ONC proposed several changes to the information blocking exception framework that have effectively split the health industry into two camps: EHR platform companies and providers on one side, and health data intermediaries, analytics companies, and patient-access advocates on the other.

Some background on how information blocking regulation works matters here. The 21st Century Cures Act created a broad prohibition on information blocking, but delegated to ONC the authority to identify “reasonable and necessary activities that do not constitute information blocking.” The exceptions are affirmative defenses. A recent Supreme Court ruling in *Cunningham v. Cornell* placed the burden of proving an exception applies on the defendant, meaning that any ambiguity in the exception framework increases litigation exposure for actors who get sued. This backdrop explains why every EHR company letter sounds alarmed even when the stated goals of the proposed changes are reasonable.

The specific exception changes ONC proposed that generated the most controversy were removing the “third party seeking modification use” condition from the Infeasibility Exception, revising or removing the “Manner Exception Exhausted” condition, restricting the Manner Exception from covering contracts of adhesion or above-market-rate agreements, and eliminating the TEFCA Manner Exception.

On removal of the third-party modification use condition, the split is clean. Companies like Innovaccer and Datavant, who are in the business of accessing EHR data on behalf of providers, supported removal. Their argument is that this condition has been used by EHR developers to block legitimate third-party integration – including write-back capabilities that make AI-generated insights actionable. When a platform tells an analytics vendor that supporting write access would require modifying their certified technology and therefore falls under the Infeasibility Exception, the exception becomes a moat rather than a safety valve. PointClickCare

Oracle, and Altera opposed removal with different levels of intensity.

PointClickCare's comment was the most direct: they argued that removing this condition combined with expanded AI/automation access would allow third-party agents to modify patient records faster than clinical staff can validate the changes they used fairly stark language about patient safety implications. Their framing that a hallucinating AI agent with mandatory write access to EHRs – which is effectively what the rule combination implies – represents a foreseeable patient scenario that ONC has not adequately analyzed.

On the Manner Exception Exhausted condition, the fight is about whether EHR vendors can require requestors to accept alternative methods of data access before declining a specific non-standard request. ONC wants to tighten the requirement essentially saying that offering one bad alternative should not count as exhausting an exception. The EHR Association and virtually every EHR company argued that the proposed changes – especially replacing “same” with “analogous” and changing “substantial number” to “any” – would introduce enormous ambiguity and compliance cost. The “analogous” standard is particularly problematic because there is no objective definition of when two APIs are analogous, which means disputes in litigation instead of getting resolved operationally. Epic was explicit that what is analogous from one technical architecture's perspective is not analogous from another's, and that this kind of subjective standard is going to generate court cases that will not be resolved well.

On the contracts of adhesion restriction, the tension is real and somewhat internally inconsistent within the proposed rule itself. ONC wants to require that agreements qualifying for the Manner Exception be at market rate, not be contracts of adhesion and not contain unconscionable terms. The problem multiple commenters identify is that ONC also requires certified API pricing to be publicly posted on a standardized basis – which looks a lot like a standardized form agreement. If standardized contracts are contracts of adhesion and therefore unavailable for the Manner Exception, while simultaneously being required for certified API pricing, those two requirements are in direct conflict. Epic noted they executed over 6,000 consultant agreements and over 300 new vendor enrollment agreements in 2025.

– at that scale, individual contract negotiation is not operationally feasible. The Association’s comment noted they could not get their member companies to agree on what parts of a technology agreement counted as related to EHI access versus incidental commercial terms, which suggests that even the regulated entities can’t reliably classify what the rule would require them to negotiate individually.

The TEFCA Manner Exception removal is the one information blocking proposal that got nearly universal support. Epic, Datavant, Premier, PointClickCare, Innovacell, NCQA, AHIP, and the EHR Association all supported removing it – though some with caveats. Epic’s comment noted that while TEFCA has scaled rapidly and now includes over 2,000 hospitals, 50,000+ clinics, and more than half a million providers, it is not yet a complete national exchange infrastructure. The argument for removal is that the exception has been used to restrict data access by pointing to TEFCA as the only acceptable pathway even when other technically capable alternatives exist, which is the opposite of what TEFCA is supposed to do. Epic suggested instead that the exception be modified rather than eliminated, specifically to give critical access hospitals and FQHCs a regulatory incentive to join TEFCA – entities that have historically not participated in national exchange frameworks at scale.

AI Transparency Gets Gutted and Nobo Is Happy About It

ONC proposed removing the so-called “model card” requirements from the Dec 2024 Support Interventions certification criterion. These were requirements introduced in HTI-1 that obligated certified EHR developers to provide source attribute transparency for predictive AI tools – information about how the model was trained, what data it uses, its performance metrics, and its intended use cases. The reason for removal is that ONC found no publicly available evidence these requirements improved patient care, and that clinicians essentially never accessed source attribute information in the workflow. Epic’s comment included a data point that in 2025, 90% of Epic-using organizations had no users who ever viewed source attributes. Other analytics showed source attribute information was accessed an average of twice per month per organization.

The problem is that removing these requirements left almost everyone uncomfortable for different reasons. Clinical organizations – ACP, MGMA, Allina Health – argue that even if clinicians are not consulting model cards at the point of care, those documents are used during implementation and procurement decisions. Removing standardized requirement means practices, especially smaller ones with limited technical staff, lose the only consistent mechanism for evaluating AI tools they are purchasing. The liability concern is real: if a practice deploys a clinical AI tool that turns out to have been trained on biased data or to perform poorly on their patient population, and there is no standardized transparency documentation, the practice carries the harm without any regulatory backstop that existed when they made the purchase decision.

Oracle's comment was nuanced – they suggested retaining the requirement that source attributes be provided to customers as product documentation, while retaining the requirement that source attributes be accessible in the EHR workflow during clinical care. This split is actually sensible and might end up being the compromise position, because it addresses ONC's legitimate observation that nobody uses model cards in real-time clinical workflow while preserving the procurement and governance function that makes them useful.

Wolters Kluwer had a different concern entirely: they raised the AI training data question, arguing that allowing autonomous AI systems to access EHI without limitation effectively enables commercial AI training on patient data without explicit authorization, creates conflicts with HIPAA's minimum necessary standard, and violate TEFCA's purpose-fidelity requirements. Their comment is probably the most legally sophisticated analysis of how the proposed AI access language interacts with existing legal frameworks, and it raises questions ONC clearly has not fully answered.

HL7 made the structural governance point: source attribute requirements and model card infrastructure are the trust rails that make safe AI integration in clinical workflows possible. Removing them before any alternative accountability framework exists does not just create a gap in the current regulatory environment – it removes the scaffolding on which a future, better framework would have been built.

The Automation and RPA Question That Could Break Everything

The most technically aggressive piece of HTI-5 is the proposal to explicitly include robotic process automation and autonomous AI systems in the definitions of “access” and “use” under the information blocking framework. The phrase ONC used – “without limitation” – became the flashpoint. What ONC intended as a forward-looking clarification that AI-enabled workflows are protected from information blocking obstruction got read by much of the industry as a mandate to treat RPA and AI agents as equivalent to human users for purposes of EHI access.

The divide here tracks roughly by business model. Datavant and Innovaccer, whose operations depend on automated workflows for processing millions of record requests, supported the definitional expansion enthusiastically. Datavant’s comment noted that processing millions of record requests annually across 80,000+ health organizations requires workflows that are automated, auditable, and high-throughput – and that automation is key to supporting these workflows in controlled ways. Innovaccer made the same point and added that as agentic AI systems increasingly mediate real-time data exchange – surfacing risk stratification alerts inside EHR workflows, transmitting updated care plans automatically – limiting the definition of “access” and “use” to human actors creates a widening enforcement gap that incumbent platforms can exploit.

PointClickCare’s response was arguably the most detailed opposition filed, and their argument deserves serious engagement rather than dismissal. Their core point is that bots overwhelm systems designed for human-speed interaction in ways that are indistinguishable in real-time from denial-of-service attacks. Distinguishing an authorized RPA workflow from a malicious scraping bot is a technically nontrivial problem that has no standard solution. Mandating that systems be open to automated access “without limitation” imposes unfunded infrastructure costs – more bandwidth, more compute, more security monitoring – on developers who priced their systems for human-speed usage. They also raised the scenario that gets ignored in most policy discussions: what happens when a third-party AI agent with write access and no

human oversight hallucinate a clinical note, medication dose, or lab result into a patient's permanent record? HIPAA audit logging requirements are not well designed to catch one-off AI errors. The downstream liability for a clinician who makes a treatment decision based on a hallucinated record entry is not adequately addressed anywhere in the proposed rule.

Epic provided three real-world examples of RPA-caused harm from their custom base without any AI involvement at all: one incident where RPA set incorrect medication doses on nearly 73,000 medications, requiring urgent remediation of 44,000 patient records; another where an RPA solution added notes to the wrong patient 90% of the time; and a third where a single RPA documentation improvement solution required an unexpected \$1M infrastructure investment to prevent system performance degradation. These are not hypothetical scenarios. They are documented operational failures at scale, and they involve simple process automation tools, not agentic AI systems. ONC is now proposing to extend equivalent access rights to

The practical resolution that most technical commenters converged on is that ONC should drop the phrase "without limitation" and instead provide specific guidance when automated access is permissible – tied to FHIR-based APIs, authenticated authorization flows, audit logging requirements, and rate-limiting controls. This framing would protect the legitimate use cases that Datavant and Innovaccer depend on while allowing developers to implement reasonable security controls without triggering information blocking liability.

What This Means for Entrepreneurs and Investors

Several commercially significant signals come through the comment letter corpus to anyone deploying capital or building companies in this space.

Lower barriers to EHR certification entry is real, but the market dynamics are complicated. Removing 34 certification criteria does reduce the compliance cost of entering the certified health IT market. The EHR Association's comment noted that the biggest single burden reduction comes from eliminating the Safety-Enhance

Design criterion, which required expensive summative usability testing. But multiple commenters noted a concern that should interest investors: lower barriers cut both ways. The certification requirements that are being removed also represented minimum quality floors that new entrants had to meet. If a new EHR company can achieve certification without demonstrating C-CDA interoperability, HIPAA-aligned security controls, or real-world testing results, the competitive pressure on established vendors may be less than it appears. The MGMA made this point explicitly – smaller practices cannot distinguish a recently certified new entrant from a system that has been certified for 15 years just by looking at the certification label.

The data intermediary thesis just got materially stronger if the information blocking changes survive. Companies in the business of accessing, normalizing, and routing EHI on behalf of providers and health plans – record retrieval, ROI, prior authorization processing, population health data aggregation – have been fighting a policy and legal battle against EHR platform owners who have used information blocking exceptions to limit their access. If ONC narrows the exceptions and extends the information blocking framework to cover automated access, that changes the leverage dynamic. The question is whether the changes that actually benefit data intermediaries – removing the third-party modification use condition, restricting adherence to the Manner Exception, codifying AI access – survive the finalization process in the face of strong EHR vendor opposition.

The FHIR API competitive layer is opening up whether the market is ready or not. This is the directional bet embedded in HTI-5 that has the most long-term significance. By removing C-CDA certification requirements and reorienting the program around FHIR APIs, ONC is effectively accelerating the moment when FHIR becomes the only game in town for interoperability. That is a large market structure shift. Every company building on EHR connectivity needs to be building on FHIR-based APIs today, not just because it is technically better but because the certification floor beneath C-CDA-based connectivity is being actively pulled away. The transition valley is real and will create short-term disruption, but the directional signal is clear and investors should be pricing it in.

The AI clinical decision support market is in a regulatory reset period with more uncertainty than opportunity in the near term. Removing model card requirements while providing no alternative transparency framework means that for a period of time – probably until HTI-6 or later guidance – the market for clinical AI tools operates without a standardized evaluation infrastructure. That is a double-edged situation. It reduces compliance cost and removes some deployment friction for vendors. It also means that health system procurement teams, who are already cautious about clinical AI liability, lose the one standardized artifact they could use in their AI governance committees. Companies selling into clinical decision support need to get ahead of this by developing strong voluntary transparency documentation, because the demand from health systems for some version of model card information is not going away just because the certification requirement was removed.

TEFCA participation is about to matter more, and the current TEFCA ecosystem reflects that. Epic alone accounts for over 63,000 directory entries in TEFCA out of what is otherwise limited broad non-Epic participation. With the TEFCA Mann Exception being removed, health systems connected through TEFCA lose one of the mechanisms they had for limiting data requests to that specific channel. This probably accelerates TEFCA as an actual interoperability channel for payers and quality measurement organizations – AHIP's comment made clear that they want to use TEFCA for HEDIS and quality measure reporting and that the current fee structure that allows providers to charge payers for that data is a problem they've not yet resolved. The emerging battle over TEFCA fee structures is one to watch closely.

On the information blocking litigation landscape, the Supreme Court's Cunniff ruling combined with ONC's proposed exception changes creates a compliance environment that disproportionately affects smaller health IT companies. Large platforms have compliance teams and legal budgets to navigate ambiguous exception frameworks. Smaller vendors – which is most of the health tech startup market – face proportionally higher compliance risk from the same ambiguity. The practical effect may be that HTI-5, framed as a deregulatory action, ends up concentrating market power in favor of incumbents through litigation risk rather than regulation.

requirement. This is something several commenters identified and something investors should model carefully when assessing the regulatory moat of health IT infrastructure plays.

The honest summary is that HTI-5 is a significant realignment of health IT regulatory architecture that creates real tailwinds for certain business models – FHIR-native infrastructure, health data intermediaries with strong API capabilities, AI companies that were already planning FHIR-first deployment – while creating substantial uncertainty for companies whose business model depends on regulatory stability, CDA-era connectivity, or clinical AI compliance documentation. The comment letters are worth reading not just for what they say about the proposed rule but for what they reveal about the strategic priorities and threat perceptions of almost every major incumbent in the space. When Epic writes 23 pages about information blocking exceptions, and when PointClickCare calls a proposed rule change potentially lethal for patients, those are signals about where the competitive pressure is accumulating and where the next several years of market structure fights will be centered.



1 Like • 1 Restack

← Previous

Next

Discussion about this post

Comments

Restacks



Write a comment...