

The Interoperability Trap: How TEFCA Design Choices Undermine Healthcare Equity for Cash-Pay Providers

MAY 09, 2025 • PAID



1



1

Share

In the evolving landscape of healthcare interoperability, there has long been a promise—a vision that, through the judicious implementation of national frameworks and trusted exchange networks, we would finally tear down the walls that have long isolated patient health data in proprietary silos. TEFCA, the Trusted Exchange Framework and Common Agreement, was heralded as a foundational step toward that vision, bringing with it the hope of seamless, modern, patient-centric access to health records. Yet, in the very early stages of its implementation, an insidious contradiction has emerged. While TEFCA promotes itself as a facilitator of equitable data access, its current operationalization actively excludes a growing segment of healthcare providers: cash-pay clinics. These are not rogue operations or fringe entities. They are state-licensed, HIPAA-compliant, often digitally forward clinical practices that, in many cases, offer superior patient engagement and continuity of care. And yet, these same clinics now find themselves shut out of the national networks they once accessed, not because they lack clinical legitimacy, but because they have opted out of the traditional insurance model.

This contradiction is not just philosophical; it is operational and profoundly damaging. Recently, I spoke with the leadership team of a primary care clinic that has structured itself around the principles of direct care. Their mission is simple: reduce administrative waste, offer transparent pricing, and prioritize the doctor-patient relationship over bureaucratic documentation. And yet, in a situation that should alarm any serious policymaker, this clinic was forced to delay care for an elderly patient because they could not receive her records in time. The hospital where she had recently been admitted, tethered to legacy workflows and suspicious of the clinic's

“non-traditional” payment model, failed to send the records electronically. Instead they were faxed—a method whose failure rate, latency, and lack of security is universally derided in the digital health space. Ironically, this clinic had been part of a Health Information Exchange just a year prior. Their credentialing had not changed. Their commitment to treatment and HIPAA compliance had not changed. What changed was the policy landscape—swept up in the aftermath of a single interoperability scandal and the reactive governance that followed.

To understand how we arrived at this impasse, one must look at the ripple effect of what occurred between a data intermediary and one of the largest electronic health record vendors in the country. The dispute began with the onboarding of certain entities to a national interoperability network under the guise of treatment-based access. While the technical definitions of treatment under HIPAA are relatively clear, the application of those definitions, when filtered through the lens of risk aversion and competitive protectionism, can be interpreted far more narrowly. A handful of onboarded organizations were found to be engaging in what appeared to be non-treatment activities, such as funneling patient data for mass tort legal actions. The data intermediary facilitating their access faced scrutiny, and rightly so. There is no place for abuse in networks designed for clinical care. But the ensuing response from dominant network stakeholders—particularly from those who control access to massive swaths of hospital and EHR data—was both sweeping and blunt. Rather than targeting the misuse with scalpel-like precision, they wielded a hammer.

The result was the application of newly stringent Standard Operating Procedure (SOPs), for vetting participants in interoperability frameworks. These SOPs, developed quickly and with limited transparency, redefined eligibility in ways that effectively equated clinical legitimacy with insurance participation. The logic, if it can be called that, was straightforward: entities participating in insurance billing were easier to track, and control. They left digital trails through claims data, could be reported on for fraud through payor relationships, and were generally part of the same ecosystem of incentives that had long governed legacy healthcare institutions. Cash-pay providers, by contrast, were harder to monitor through traditional mechanisms and thus, in the eyes of SOP authors, posed a risk.

It is here that the fundamental philosophical flaw reveals itself. Healthcare interoperability was never meant to be a compliance mechanism for insurance infrastructure. It was not conceived as a tool to entrench existing business models. It was designed to empower patients and providers—regardless of payment modalities—to share, receive, and act upon health information in a way that improves outcomes and reduces administrative friction. To now see this vision narrowed by gatekeeping criteria that are entirely financial in nature is to witness the perversion of the original goal.

The reality is that cash-pay providers are not inherently riskier. In fact, many direct care clinics offer longer visits, lower patient panels, and tighter operational controls precisely because they are not beholden to the throughput demands of fee-for-service billing. These providers are often innovators in care delivery, leveraging digital tools, asynchronous communication, and personalized medicine models to better serve underserved communities. Their exclusion from national networks is not only a disservice to patients—it is an institutional betrayal of the principles of clinical neutrality and access equity.

The problem is compounded by the opaque nature of how these SOPs were created and who gets to enforce them. There is no public docket of stakeholders, no opportunity for comment or redress, no clear appeals process. The entities that maintain technical infrastructure for interoperability—be they Qualified Health Information Networks under TEFCA, or older frameworks like Carequality and CommonWell—are making access decisions that affect hundreds of thousands of providers and millions of patients. Yet these decisions are happening behind closed doors, with input from a limited circle of enterprise-focused participants. This lack of transparency breeds distrust and creates an uneven playing field where influence and integrity determine access.

One of the great ironies in all of this is that the very actors the SOPs were designed to exclude—fraudulent telehealth operators, shell clinics, and data brokers posing as providers—can still often gain access simply by submitting the right insurance credentials. They may bill minimal amounts, use third-party credentialing services, and operate out of virtual addresses, but so long as they exist on paper as part of the

traditional insurance system, they can slip through. Meanwhile, legitimate clinics with brick-and-mortar locations, robust clinical teams, and full HIPAA compliance are denied access because they have chosen to remove insurance from their operations. This is not risk management. It is risk misallocation.

A more sophisticated and fair approach would acknowledge the complexity of medical care delivery and build fraud detection mechanisms that are behavior-based rather than status-based. Instead of assuming risk based on the presence or absence of insurance billing, networks should be analyzing query patterns, attestation quality, and clinical documentation. They should deploy anomaly detection to flag unusual access behavior, implement graduated penalties for non-compliance, and create processes for remediation. None of this requires the exclusion of an entire class of providers. It simply requires better tooling, better governance, and a willingness to evolve past legacy heuristics.

At a deeper level, this debate touches on the role of interoperability in a system that is increasingly being shaped by alternative models of care. The rise of direct primary care, virtual-first care, behavioral health startups, and chronic care platforms has introduced new operational structures that do not fit neatly into the mold of twentieth-century healthcare institutions. If interoperability is to be meaningful in this new landscape, it must adapt to reflect the diversity of care models now serving patients. TEFCA's promise lies in its potential to standardize data exchange across this diversity, not in its capacity to reinforce a narrow definition of provider legitimacy.

The policy path forward is clear, though not politically easy. First, there must be immediate effort to create an inclusion pathway for cash-pay providers that is based on clinical credentials and compliance obligations, not financial arrangements. A provider with state licensure, an NPI, and demonstrable HIPAA compliance should be eligible to participate in treatment-based data exchange under TEFCA. This does not open the floodgates to abuse; it simply realigns eligibility with clinical reality. Second, the governance structures that define SOPs and operational policies for interoperability networks must be democratized. Whether through public comment periods, advisory boards with diverse representation, or transparent publication

decision-making processes, the time has come to open the black box. Third, fraud prevention must shift from gatekeeping to intelligent monitoring. Just as banks shut out all small businesses for fear of money laundering but instead monitor transactions for suspicious activity, interoperability networks must do the same data access.

TEFCA is still young. Its architecture is flexible enough to accommodate course corrections. But if it continues down its current trajectory, it risks becoming an enforcer of status quo healthcare power dynamics rather than a catalyst for patient empowerment and innovation. The goal should never have been to protect institutions from innovation. It should have been to protect patients from information asymmetry, administrative opacity, and the clinical dangers of fragmented care. Cash-pay providers are not a fringe group—they are a growing segment of the healthcare delivery landscape, one that is often closer to the patient, more responsive to clinician nuance, and more willing to embrace new care paradigms. Excluding them from national health data conversation is not just policy malpractice. It is a strategic failure.

In the coming years, as TEFCA expands into use cases like Individual Access Services and Public Health Reporting, the foundational decisions made today will reverberate across every layer of our health system. If we cannot get the basic use case of treatment right—if we cannot ensure that any legitimate provider treating a patient can access the data needed to do so safely—then the rest is merely window dress. It is incumbent upon those shaping these frameworks to revisit their assumptions, broaden their perspectives, and return to the first principles of clinical care. Data belongs to the patient. Providers serve the patient. And interoperability exists to support that sacred relationship, not to constrain it through outdated proxies for legitimacy.

The exclusion of cash-pay providers from national interoperability frameworks in TEFCA represents a form of systemic epistemic blindness, one that fails to comprehend the trajectory of healthcare delivery in the United States. What once have been considered an outlier model—where patients pay directly for care, and providers structure their business models around clinical value rather than insu-

billing—has now become a foundational pillar of primary and specialty care innovation. These clinics are often early adopters of new health technologies, electronic health record systems built for usability rather than compliance check and care models that emphasize continuity over volume. Yet TEFCA's current implementation, rooted in the protective instincts of large institutional players, to recognize or accommodate this evolution.

What exacerbates the problem is that these exclusions are not based on clear evidence of risk but on an assumption that insurance billing serves as a proxy for legitimacy. This logic is dangerously flawed. Insurance billing does not guarantee that a provider is delivering high-quality care. It does not ensure ethical data usage, nor does it provide assurance against fraud. In fact, some of the most egregious instances of healthcare fraud in recent history have occurred within the insured system, where phantom billing, upcoding, and patient data manipulation are not just possible—they are widespread. By contrast, many direct care providers operate in a transparent and accountable manner precisely because their patients are their only source of revenue and reputation.

The deeper issue is that the current framework for TEFCA, and by extension the broader interoperability regime, has not been built with adaptive governance in mind. Governance in health IT tends to congeal quickly. Once a policy is set, particularly one as foundational as eligibility criteria for network participation, it becomes difficult to unwind. The stakeholders involved—large EHR vendors, hospital systems, and government agencies—often converge around risk-averse practices that favor incumbency and bureaucratic inertia. This is not because they are malicious. It is because they are accountable to large systems, complex risk portfolios, and compliance burdens that resist ambiguity. But in doing so, they build policy environments that are brittle, exclusionary, and ultimately unfit for a pluralistic healthcare economy.

To rectify this, a fundamental rethinking of interoperability governance is required—one that moves away from static credential-based exclusion toward dynamic, behavior-based inclusion. Imagine an interoperability framework where access is not determined by insurance contracts or billing codes but by adherence to transparent

clinical standards and observable behavior. Such a framework would allow for continuous monitoring of data access patterns, flagging anomalous behaviors that deviate from expected clinical workflows.

This model of behavior-based access is not hypothetical. It draws from principles already in use in other sectors where sensitive data exchange occurs at scale. Consider the financial services industry, which processes trillions of dollars in daily transactions. Access to core banking APIs, international wire transfers, and real-time payment systems is not granted solely based on whether a business fits a traditional mold. It is granted through a blend of credential validation, transaction monitoring, and adaptive risk scoring. The notion that a provider must prove legitimacy by aligning with a legacy business model would be laughable in banking or logistics; healthcare continues to conflate operational orthodoxy with trustworthiness.

If TEFCA is to fulfill its mission, it must integrate a similar logic. The first step in doing so is to formalize the concept of continuous vetting. Rather than relying on a one-time approval process or manual SOP-based screening, participation in national data networks should be governed by persistent verification, with real-time monitoring of data use behaviors. This would allow the system to distinguish between providers who make legitimate treatment-based queries and actors who seek to exploit the network for purposes inconsistent with patient care. The metrics for such monitoring already exist—query frequency, target diversity, response utilization, and even alignment between query type and specialty can all serve as inputs into a risk profile. When unusual patterns are detected, a tiered escalation protocol can guide review, rather than triggering a punitive expulsion that penalizes providers for systemic design flaws.

A continuous vetting framework also allows for flexibility across care models. A small-volume mental health practice in a rural town will not behave like a high-acuity urgent care chain in an urban center. Nor should it have to. The system must be intelligent enough to understand these nuances. Risk is not a monolith, and neither is legitimacy. By embracing variability rather than resisting it, TEFCA could create a framework that is resilient, adaptive, and inclusive—all without sacrificing security or compliance.

Yet none of this can occur unless the governance structures that oversee TEFCAs are made transparent and participatory. Currently, a small number of organizations play an outsized role in determining who gains access and under what conditions. These entities, while technically sophisticated and operationally necessary, do not represent the full spectrum of healthcare delivery. They are disproportionately staffed and influenced by stakeholders whose interests are aligned with institutional scale, reimbursement infrastructure, and network control. The voices of small practices, virtual care providers, and cash-pay clinics are effectively absent from the rooms where these decisions are made.

This is a governance failure, not a technological one. TEFCA's governing bodies must include seats for cash-pay clinics, virtual-first care providers, rural practitioners, and others operating outside the traditional healthcare establishment. These stakeholders must not be token voices, but equal participants in the deliberative process that shapes national data exchange policy.

Furthermore, a formal appeals process should be established—one that allows data providers to present evidence of legitimacy, intent, and compliance. This process should be handled by a neutral entity with the authority to override or modify decisions, ensuring that access is not arbitrarily denied. The appeals process should also include published decisions with redacted details, to establish precedent and ensure accountability. Without such mechanisms, the onboarding process becomes opaque, bureaucratic, and susceptible to institutional bias.

Another key component is building a technical monitoring infrastructure that can support proportional oversight. Instead of assuming bad faith and designing policies to exclude, TEFCAs must assume good faith and verify it through intelligent audit mechanisms. These can include real-time query dashboards, machine learning models that detect anomalies, and collaborative review teams that include clinicians, data scientists, and compliance officers. When misuse is detected, sanctions should be proportionate—ranging from queries suspension to access revocation—but always transparent, appealable, and rooted in clinical and ethical analysis.

In the end, the health of any interoperability framework is measured not by the rigidity of its gates, but by the integrity of its flows. TEFCA was built to facilitate trust, but trust is not a static credential. It is an ongoing process of transparency, participation, and shared responsibility. When that process excludes entire categories of providers for failing to adhere to legacy economic models, it betrays the very concept it was meant to uphold.

We are standing at a critical juncture. The policies we implement today will define the shape of our data ecosystems for decades to come. Will they be open, adaptive, and patient-centered? Or will they become yet another mechanism through which entrenched power structures control access to the tools of modern care?

The future of healthcare will not be uniformly institutional. It will be pluralistic, messy, dynamic, and decentralized. Our interoperability systems must reflect that reality. If we fail to adapt, we will have built the infrastructure of yesterday for the patients of tomorrow. And in doing so, we will have squandered the rarest of opportunities: the chance to build a health data system worthy of the trust we so claim to defend.



1 Like • 1 Restack

← Previous

Next →

Discussion about this post

Comments

Restacks



Write a comment...