Reducing Fraud and Improving Payment Integrity in Healthcare Through the Use of AI

May 2025

HL7 International

About HL7 International and its Work Related to AI

Founded in 1987, Health Level Seven International (HL7) is a not-for-profit, ANSI-accredited standards development organization with a long history of forging consensus on and driving adoption of standards for the exchange, integration, and retrieval of health information to support improvements in health and healthcare.

In 2024, HL7 launched an effort focused on artificial intelligence[1] led by HL7 CEO Charles Jaffe, MD, PhD. In its first year, HL7 launched three strategic initiatives, the first of which is focused on countering fraud. This report, "Reducing Fraud and Improving Payment Integrity in Healthcare Through the Use of AI" is the first key output of this strategic initiative—which is led by co-chairs Janet Marchibroda and J. Marc Overhage, MD, PhD in collaboration with David Bray, PhD, MSPH of the Stimson Center.

The report cover was generated using AI, specifically FLUX Pro.

For more information about HL7, visit www.hl7.org

[1] For details see: https://info.hl7.org/ai

Forward and Letter From the CEO of HL7

The United States spent more than $4.9 trillion on healthcare in 2023, accounting for 17.6 percent of the nation's gross domestic product (GDP)[2]. Healthcare costs are continuing to increase, with growth of 7.5 percent over the last year alone[3] .

In addressing rising healthcare costs, every participant in the health care continuum must focus on those costs attributable to fraud, waste, and abuse as they are significant. As such, these avoidable systematic elements ultimately cost the healthcare system more than $900 billion per year. Several other nations across the globe are experiencing similar cost challenges associated with providing healthcare.

The use of artificial intelligence (AI) promises to significantly improve payment integrity methods which help ensure accurate claims processing and payment and reduce fraud, waste, and abuse.

Recognizing the important role of standards, electronic data capture for, and integration of the results of these AI-based tools into workflows, HL7 convened a series of discussions with healthcare payers, providers, and technology experts to assess current challenges in the use of AI for payment integrity and fraud detection and prevention. More importantly, they summarized the actions needed to fully realize the benefits of AI, with the goal of reducing overall healthcare costs.

This report summarizes the current landscape, including interoperability and data access barriers and AI implementation challenges, as well as opportunities for enhancing the use of AI for improving payment integrity and addressing fraud.

For the readers of this report, I recommend you consider the importance and strong value of a neutral convener of industry and public sector efforts in working to

[2] Centers for Medicare and Medicaid Services. National Health Expenditure Accounts. https://www.cms.gov/data-research/statistics-trends-and-reports/national-health-expenditure-data/historical

[3] Ibid.

improve healthcare, to include reducing fraud and improving payment integrity, through next-generation AI technologies.

Key recommendations from the report include

- Developing standards for explainable AI in healthcare with specific transparency requirements and bias mitigation protocols.
- Creating standards and infrastructure for trust and verification of AI-generated results.
- Establishing frameworks for human-in-the-loop validation that balance automation with clinical expertise.
- Implementing pilot programs focused on provider-payer collaboration that create mutually beneficial scenarios for all stakeholders.

HL7 is a neutral body with a decades-long history of gaining consensus on and supporting the implementation of standards and best practices for interoperability and electronic information sharing. These are essential components of any AI solution in healthcare. As such, HL7 stands ready to support the rapidly evolving field of healthcare and the important role that AI plays in improving health and healthcare in the United States and across the globe.

I commend the work of the Co-Chairs of the Strategic Initiative on Use of AI to Counter Fraud and Improve Payment Integrity (Strategic Initiative on AI to Counter Fraud), J. Marc Overhage, MD, PhD and Janet M. Marchibroda, as well as David Bray, PhD, MSPH as Chair of the Accelerator and Distinguished Fellow at the Stimson Center, for their collective work and significant contributions in convening these discussions, synthesizing the inputs, and producing this report. In addition, I thank the multiple individuals who provided their expert input, non-attributed thoughts, and feedback to this report both through formal interviews and information conversations with the members of the Strategic Initiative on AI to Counter Fraud.

Together, we can harness the transformative power of AI to create a more efficient, equitable healthcare system that benefits patients, providers, and payers alike. I invite you to join us on this critical journey as we work to reduce costs, improve care, and build a healthier future for all.

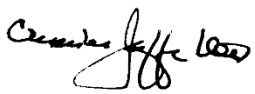

Charles Jaffe, MD, PhD
CEO, HL7 International

Executive Summary

Health Level Seven International (HL7) convened a series of discussions with healthcare payers, providers, and technology experts to explore how artificial intelligence (AI) can address fraud and improve payment integrity. This report synthesizes key findings and recommendations from these discussions, focusing on interoperability standards, data governance frameworks, and implementation strategies that balance technical capabilities with practical implementation considerations.

Healthcare costs totaled $4.9 billion in the United States in 2023, increasing by an alarming 7.9 percent over the previous year[4]. Moreover, current estimates indicate that more than $900 billion in the United States is lost because of fraud, waste, and abuse. Innovative strategies are needed not only to improve the delivery of care and the management of health, but also to assure that payments are accurate, that healthcare providers are reimbursed correctly for care delivery, and that fraud, waste, and abuse are both detected and prevented.

As part of the Strategic Initiative on Use of AI to Counter Fraud and Improve Payment Integrity (more briefly: Strategic Initiative on AI to Counter Fraud), HL7 convened roundtable discussions with multiple people to seek their expert input, non-attributed thoughts, and feedback to inform this report both through formal interviews and information conversations. Our discussions included observations from the roundtable participants that they had seen multiple instances of healthcare claims coded and processed incorrectly – not necessarily because of fraud, but often because of difficulties in identifying the right billing codes.

In addition, participants highlighted that new value-based care models that promise to significantly improve quality and reduce costs will introduce new complexities related to calculation of payment. Roundtable participant discussions also included observations that current claims review processes cost hundreds of dollars per claim. Cumulatively, these challenges can create a financial burden and an environment where administrative errors as well as intentional fraud can persist undetected and unresolved. Yet hope exists to tackle these challenging issues, to include solutions that can not only process valid claims more quickly but also detect and prevent anomalous and fraudulent ones faster too.

Participants identified significant opportunities to leverage next-generation AI approaches for fraud detection and prevention while simultaneously improving payment accuracy and timeliness, including those related to new value-based healthcare models,

[4] Centers for Medicare and Medicaid Services. National Health Expenditure Accounts. https://www.cms.gov/data-research/statistics-trends-and-reports/national-health-expenditure-data/historical

to significantly reduce billing and payment processing costs. Based on all the inputs received from multiple participants, **the three key findings of our report include:**

1. **AI has significant potential to transform healthcare fraud detection and payment integrity, while also improving the transition to value-based care**.
2. **Successful implementation requires thoughtful standards development and frameworks that address explainability, transparency, data governance, and validation**.
3. **Pilot programs focused on provider-payer collaboration represent a promising path forward**.

Participants further emphasized that standards development should lead to actionable outcomes rather than theoretical frameworks, with particular attention to facilitating faster valid claims processing while effectively identifying payment errors or fraudulent activities early in the process, prior to payment.

Immediate Recommendations	Envisioned Outcomes
Commence working groups to develop standards or frameworks for explainable AI, transparency, data segmentation, and validation.	Trusted AI tools are adopted to improve payment integrity and identify potential fraud while simultaneously streamlining valid claims processing.
Initiate planning for pilot programs focused on provider-payer collaboration to improve payment integrity.	Providers and payers collaborate effectively around payment, supported by interoperable systems and standards, to reduce administrative burden and improve payment accuracy.
Engage with federal and state agencies to ensure alignment with existing and emerging regulatory and payment approaches.	AI models and related standards reflect and inform evolving regulatory requirements and payment models, including value-based care.

J. Marc Overhage, MD, PhD
Co-Chair, Strategic Initiative on Use of AI to Counter Fraud and Improve Payment Integrity

Janet M. Marchibroda
Co-Chair, Strategic Initiative on Use of AI to Counter Fraud and Improve Payment Integrity

David Bray, PhD, MSPH
Chair of the Accelerator and Distinguished Fellow
Henry S. Stimson Center

Table of Contents

1. Introduction and Background

1.1 Context and Need

The U.S. healthcare system continues to face significant challenges related to fraud, waste, and inefficiency. One study estimates that 3 percent of total healthcare expenditures associated with the U.S. healthcare system are lost to healthcare fraud[5]. Another study suggests that clinical waste accounts for an additional 5.4 percent to 15.7 percent of healthcare spending[6], resulting in anywhere from $412 billion to $916 billion lost due to fraud, waste, and abuse.

In addition, the healthcare system is increasingly shifting to value-based care models to improve quality, cost, and health outcomes, with an estimated 60 percent of total care delivery reimbursement representing such models[7]. These models bring a considerable increase in the complexity associated with calculations of payment, introducing the potential for an increase in new payment errors.

The use of AI promises to significantly improve fraud detection and prevention and payment integrity models, including those that pertain to value-based care, resulting in improvements in administrative efficiency, coding accuracy, healthcare provider (and patient) experiences, and ultimately medical cost savings.

As healthcare organizations increasingly adopt these promising AI technologies, there is both opportunity and risk in how these tools are deployed to address these challenges and realize these improvements. The current friction around accessing medical records creates significant barriers to effective payment integrity and fraud detection and prevention-related activities, with the scarcity of qualified workforce to review records further limiting the discovery of errors and potential fraud. AI technologies offer the potential to scale these review processes dramatically while reducing costs, but require thoughtful implementation guided by appropriate standards and governance frameworks.

The HL7 discussions were initiated to explore the intersection of AI, interoperability standards, and healthcare fraud detection and prevention, as well as payment integrity. These discussions, which occurred between January and April 2025 with participation

---

[5] National Health Care Anti-Fraud Association. The Challenge of Health Care Fraud. https://www.nhcaa.org/tools-insights/about-health-care-fraud/the-challenge-of-health-care-fraud/

[6] Health Affairs Research Brief. The Role of Clinical Waste in Excess U.S. Health Spending. June 9, 2022, Updated October 6, 2022. https://www.healthaffairs.org/content/briefs/role-administrative-waste-excess-us-health-spending

[7] Health Care Payment Learning and Action Network. *2022 APM Measurement Effort*. November 9, 2022. https://hcp-lan.org/workproducts/apm-methodology-2022.pdf

from a diverse group of stakeholders, including healthcare payers, providers, and technology experts, informed our understanding of current challenges and development of recommended solutions, which are detailed in this report.

1.2 Scope and Objectives of this Report

The discussions convened by the Strategic Initiative on AI to Counter Fraud focused on three primary objectives:

1. **Exploring how AI impacts risk assessment and helps address fraud and improves payment integrity in healthcare**: The discussions examined both the potential benefits of AI in detecting fraud and payment errors, as well as the risks associated with AI implementation, including concerns about bias, transparency, and regulatory compliance. Participants explored how AI could help identify not just isolated instances of fraud but also connect seemingly unrelated cases to uncover larger patterns of organized fraud.
2. **Identifying opportunities to increase healthcare efficiency and effectiveness through AI implementation**: Beyond fraud detection, the discussions explored how AI could improve overall healthcare efficiency by streamlining valid claim processing and reducing administrative burden, while also supporting the transition to value-based care models. Participants emphasized that effective solutions must create value for all stakeholders in the healthcare ecosystem, not just address fraud detection in isolation.
3. **Developing recommendations for standards and best practices to support these goals**: The HL7 discussion focused on developing practical, implementable recommendations for standards and best practices that could guide the responsible deployment of AI in healthcare fraud prevention and detection, as well as payment integrity. These recommendations encompass technical standards, governance frameworks, and implementation strategies, with particular attention to ensuring that standards development leads to actionable outcomes rather than just theoretical frameworks.

The scope encompassed both technical considerations (such as data interoperability and AI model validation) and practical implementation challenges (including workflow integration and stakeholder alignment). The discussions also addressed policy and regulatory considerations, recognizing that effective solutions would require alignment with existing regulatory frameworks and potential evolution of these frameworks to address emerging AI capabilities.

2. Current Landscape and Challenges

2.1 The Scale of Payment Inaccuracy and Fraud

Discussions revealed concerns surrounding the number of claims coded or paid incorrectly, highlighting the magnitude of payment accuracy challenges in healthcare. While the exact proportion of healthcare claims initially paid incorrectly can vary based on the payer and specific circumstances, data suggests that error rates typically range from around 6-to-20-percent[8]. These payment errors can result from several factors, including administrative mistakes, insufficient documentation, and issues with prior authorization. One participant shared an anecdote about how they and their spouse received identical blood tests during annual physical exams, yet the provider coded these tests completely differently on their respective claims, illustrating the inconsistency in current billing practices.

Current claims review processes are expensive, costing hundreds of dollars per claim, making comprehensive review economically infeasible. This creates an environment where both unintentional errors and deliberate fraud can persist undetected. The inflated cost of review is driven by several factors, including the need for specialized clinical expertise to interpret medical records, the time-intensive nature of manual review processes, and the friction associated with accessing and exchanging medical records across different healthcare systems.

Participants noted that the line between fraud and payment integrity issues is often thin, frequently boiling down to intent rather than the outcome. Payment integrity issues may stem from technical glitches, policy insufficiencies, or contradictory guidance , particularly in rapidly evolving areas like telehealth. This complexity makes it challenging to develop systems that can effectively distinguish between intentional fraud and unintentional errors, requiring nuanced approaches that consider context and patterns rather than isolated incidents.

The discussions also highlighted recent developments in the healthcare fraud landscape, including ongoing investigations into major healthcare organizations, emphasizing the continued relevance and importance of effectively addressing payment integrity and fraud in healthcare payment systems at both the individual and organizational levels.

[8] https://www.kff.org/private-insurance/issue-brief/claims-denials-and-appeals-in-aca-marketplace-plans-in-2023/?utm_source=chatgpt.com

2.2 Interoperability and Data Access Barriers

Access to medical records while protecting patient privacy remains a significant challenge in the healthcare system. Participants highlighted how the lack of standardized approaches to medical record sharing creates inefficiencies and impedes effective payment integrity and fraud prevention and detection processes. The scarcity of qualified workforce to review records further limits the discovery of errors and potential fraud.

While emerging AI technologies could help scale record review processes and reduce costs significantly, these benefits cannot be fully realized without addressing underlying interoperability challenges. The need for human-in-the-loop validation was raised by participants in the discussions, specifically in the context of AI in healthcare.

Participants emphasized that the ability to access and analyze medical records is crucial for validating claims and identifying discrepancies, yet current systems make this process unnecessarily cumbersome and expensive. Participants also discussed how information sharing extends beyond traditional healthcare systems that may then be captured by AI systems.

Healthcare data requires clearer categorization systems based on sensitivity levels to enable appropriate sharing, to include varying regulations (e.g., by state) and policies, while also maintaining privacy protections. Participants noted that the current binary approach to healthcare data privacy does not adequately address the nuanced nature of several types of healthcare information, with some data being more sensitive than others (e.g., behavioral health, HIV status, or addiction treatment information).

This creates additional challenges for data governance and privacy protection, requiring frameworks that can address both traditional healthcare data exchange and newer forms of health information sharing in the digital ecosystem.

2.3 AI Implementation Challenges

Participants identified several challenges specific to AI implementation in healthcare fraud detection:

1. **Model Validation and Trust**: The rapid evolution of AI technologies has created uncertainty about how to properly validate models, particularly when they may be used for fraud detection that could impact provider reputations and livelihoods. Participants emphasized the need for standardized approaches to model validation, including the use of benchmark datasets that represent the diversity of healthcare delivery contexts and clear metrics for evaluating model performance.

2. **Hallucination and Bias**: Concerns about AI "hallucination" (generating plausible but incorrect information) and bias were frequently mentioned, with participants suggesting the need for clear metrics for acceptable levels of error (suggested at less than two-percent). Participants emphasized that AI systems used for healthcare fraud detection must be held to particularly exacting standards of accuracy and fairness, given the potential consequences of false positives or biased outcomes.

3. **Human-in-the-Loop Requirements**: There was strong consensus that AI systems for fraud detection must include human expert validation, particularly for clinically complex decisions. Participants emphasized that AI should be viewed as a tool to augment human expertise rather than replace it entirely, with clear frameworks for determining when and how human review should be incorporated into AI-assisted processes.

4. **Regulatory Alignment**: Participants noted challenges in aligning AI implementations with existing regulatory frameworks, particularly when those frameworks were developed before current AI capabilities existed. The discussions highlighted the need for both evolution of regulatory frameworks to address emerging AI capabilities and development of AI systems that can operate effectively within existing regulatory constraints.

Participants also discussed the challenge of ensuring that AI systems can adapt to evolving fraud patterns while maintaining regulatory compliance. This requires approaches to continuous model updating and version control that balance the need for adaptation with the requirements for validation and documentation in healthcare systems.

3. Opportunities and Emerging Solutions

3.1 Balancing Fraud Detection with Healthcare Effectiveness

The discussions highlighted the potential for a more balanced approach to healthcare payment systems. Rather than focusing exclusively on detecting fraudulent billing practices, participants suggested developing standards that could also provide data-driven guidance to help providers accurately and efficiently capture information to code and/or calculate payments owed and bill for services in compliance with payment contracts and regulatory requirements. This dual approach would create a more balanced system that protects payers from abuse while ensuring providers have the tools to capture appropriate reimbursements.

Participants emphasized that effective solutions must create value for all stakeholders in the healthcare ecosystem, not just to address fraud detection in isolation. By developing standards and systems that simultaneously streamline valid claim processing while identifying potential fraud, the healthcare system could reduce administrative burden while improving payment accuracy. This approach recognizes that most healthcare providers are not engaging in fraudulent practices and would benefit from systems that facilitate faster, more accurate payment for legitimate services.

Participants also emphasized that this approach would be particularly valuable in supporting the transition from fee-for-service to value-based care models. By developing standards that support more accurate and efficient capture of data to support outcomes-based payment, the healthcare system could ease the transition to payment models that align incentives around quality and outcomes rather than volume of services.

The discussions also highlighted the importance of considering the unique challenges faced by several types of healthcare providers, particularly those in rural or underserved areas. These providers may have limited resources for navigating complex billing systems and could benefit significantly from standards and tools that help them establish appropriate billing practices while reducing administrative burden.

3.2 AI-Assisted Medical Record Review

Emerging AI technologies offer significant potential to transform medical record review processes. Participants discussed how AI could help scale review processes that are currently prohibitively expensive, potentially reducing costs from hundreds of dollars per claim to approximately one dollar per record. This dramatic cost reduction could enable

more comprehensive review of claims, leading to better detection of both unintentional errors and deliberate fraud.

AI systems can interpret free-form prose in medical records, addressing one of the key challenges in current review processes. This capability allows for more efficient extraction of relevant information from clinical documentation, enabling faster and more accurate validation of claims. However, participants emphasized that the effectiveness of these systems depends on creating appropriate incentives for information to flow through healthcare systems, addressing the underlying interoperability challenges.

HL7 Standards Support for Humans in the Loop

Both private and public sector organizations can employ existing HL7 FHIR resources like 'Task', 'AuditEvent' and 'DetectedIssue' which can be leveraged to define standardized workflows for human-in-the-loop validation.

Extensions or profiles to these resources could more clearly specify roles, responsibilities, audit requirements, and decision accountability within clinical or administrative contexts.

Participants also discussed the potential for AI to help identify patterns across seemingly isolated cases, potentially uncovering larger schemes of organized fraud. This capability could be particularly valuable for addressing sophisticated fraud operations that might otherwise escape detection through traditional review processes focused on individual claims.

However, participants emphasized that successful implementation would require:

1. **Standards for AI verification and validation** that combine generative capabilities with rule-based verification systems. Participants discussed the need for "checks and balances" systems that ensure AI-generated analyses are accurate and reliable, particularly when used for fraud detection that could have significant consequences for providers.
2. **Frameworks for human-in-the-loop validation,** especially for clinical decisions. These frameworks should clearly define when and how human expertise should be incorporated into AI-assisted processes, ensuring that complex clinical judgments are not left entirely to automated systems.
3. **Clear requirements for model transparency and validation** using standardized datasets. Participants emphasized the importance of being able to explain and justify AI-generated analyses, particularly when these analyses might be used in fraud investigations or payment disputes.

3.3 Provider-Payer Collaboration

A recurring theme in the discussions was the need for improved collaboration between providers and payers. Participants suggested that AI-enabled systems could create "win-win" scenarios where both providers and payers would benefit from more accurate and efficient claims processing. This collaborative approach represents a shift from the at times adversarial relationship between providers and payers to a more partnership-oriented model focused on shared goals of payment accuracy and timeliness of payment.

One participant proposed a pilot program focusing on collaboration between a hospital and a payer where medical records would flow freely, enabling a lower-friction way to get providers paid more accurately upfront. This approach could simultaneously reduce administrative burden for providers while helping payers identify potentially fraudulent claims more effectively. By focusing on specific high-volume claim types where current processes create friction, such pilots could demonstrate measurable improvements in both payment accuracy and administrative efficiency.

Participants emphasized that successful collaboration would require alignment of incentives between providers and payers, with clear benefits for both parties. For providers, these benefits would include faster more accurate payment and reduced administrative burden. For payers, benefits would include better fraud detection, improved payment accuracy, and potentially lower administrative costs associated with claims processing and review.

The discussions also highlighted the importance of engaging frontline stakeholders in the development of collaborative approaches, ensuring that solutions address real-world challenges and can be effectively implemented in diverse healthcare settings. This engagement would help ensure that standards and systems developed by HL7 for example, would lead to actionable outcomes rather than just theoretical frameworks.

4. Recommendations for Standards Development

4.1 Explainable AI and Transparency Standards

HL7 Standards Support for AI Results and Traceability

Organizations can employ existing HL7 FHIR resources such as 'Provenance', 'AuditEvent', 'DiagnosticReport', and 'Observation' to encapsulate AI outputs with full traceability.

HL7 Implementation Guides can profile these resources, while CQL/FHIRPath expressions in the Clinical Reasoning module ('Measure''MeasureReport') can formalize validation rules.

The discussions yielded staunch support for developing standards related to explainable AI in healthcare, particularly for applications involving fraud detection. Participants emphasized that AI systems used for healthcare fraud detection must be transparent and explainable, given the potential consequences of false positives or biased outcomes.

Specific recommendations regarding AI results and traceability include:

1. **Transparency Requirements**: Develop standards that define minimum requirements for transparency in AI systems used for healthcare fraud detection, including documentation of training data sources, model architecture, and performance metrics. These standards should ensure that stakeholders can understand how AI systems reach their conclusions, particularly when these conclusions might impact provider reputations or payment decisions. Participants suggested that transparency requirements should be tailored to different stakeholders' needs, with distinct levels of detail appropriate for technical experts, clinicians, administrators, and regulatory bodies.
2. **Bias Mitigation Protocols**: Establish standardized approaches to identifying, measuring, and mitigating bias in AI systems, with particular attention to ensuring that fraud detection does not disproportionately impact specific provider types or patient populations. These protocols should include requirements for diverse training data, regular bias audits, and corrective measures when bias is detected. Participants emphasized that bias mitigation is particularly important in healthcare fraud detection, where false positives could have significant consequences for providers and patients.
3. **Explanation Frameworks**: Create standardized frameworks for generating explanations of AI decisions that are meaningful to different stakeholders, including clinicians, administrators, and regulatory bodies. These frameworks should ensure that explanations are both technically accurate and practically

useful, enabling stakeholders to understand and potentially challenge AI-generated analyses. Participants suggested that explanation frameworks should be tailored to different use cases, with distinct levels of detail appropriate for routine claim processing versus potential fraud investigations.

Participants also emphasized the importance of establishing clear metrics for acceptable levels of error in AI systems used for healthcare fraud detection, with some suggesting that error rates should be kept below two percent. These metrics should address both false positives (incorrectly flagging legitimate claims as potentially fraudulent) and false negatives (failing to identify fraudulent claims), recognizing the different consequences associated with each type of error.

4.2 Data Governance and Trust Infrastructure

Participants emphasized the need for robust data governance frameworks and trust infrastructure to support AI implementation in healthcare fraud detection:

1. **Data Categorization Systems**: Develop standards for categorizing healthcare data based on sensitivity levels to enable appropriate sharing while maintaining privacy protections. Participants noted that current HIPAA regulations take a relatively binary approach to healthcare privacy, which does not adequately address the nuanced nature of several types of healthcare information. A more granular categorization system would enable more appropriate sharing of less sensitive information while maintaining strict protection for extremely sensitive data (such as behavioral health information, HIV status, or addiction treatment).
2. **Trust and Verification Infrastructure**: Create infrastructure standards to ensure trust and verification of AI-generated healthcare data and claims, including approaches to digital provenance and chain of custody. These standards should enable stakeholders to verify the source and integrity of data used in AI systems, as well as the information outputs from these systems. Participants emphasized the importance of establishing clear mechanisms for validating AI-generated analyses, particularly when these analyses might be used in fraud investigations or payment disputes.
3. **Standardization at the National and State Levels**: Address both national-level and state-level standardization efforts, potentially working with state and regional CMS initiatives on data cooperatives and data trusts. Participants noted that healthcare regulation occurs at both state and federal levels, creating challenges for standardization efforts. Effective standards development would need to consider this multi-level regulatory environment and develop approaches that can work across different districts.

Participants also discussed the challenges associated with information sharing beyond traditional healthcare systems, with patients often sharing health information on consumer-facing applications that may then be captured by AI systems. This creates additional challenges for data governance and privacy protection, requiring frameworks that can address both traditional healthcare data exchange and newer forms of health information sharing in the digital ecosystem.

HL7 Standards Support for Explainable AI Artifacts

For explainability, a FHIR 'Evidence' resource, paired with the emerging 'Justification' element pattern, can serialize the reasoning trace for an AI decision. Linking the AI output 'Observation' to an 'EvidenceVariable' instance clarifies which data features contributed to the conclusion.

For AI error metrics, existing HL7 FHIR 'MeasureReport' and 'Observation' can be leveraged to introduce standardized error reporting formats. HL7 could explicitly define extensions or value sets for reporting AI model performance metrics (accuracy, specificity, sensitivity, error rates) within these resources, ensuring consistent and transparent reporting across AI implementations in healthcare fraud detection.

4.3 Model Updating and Regulatory Compliance

The discussions highlighted the dynamic nature of healthcare fraud and the need for AI systems to adapt while maintaining regulatory compliance:

1. **Continuous Model Updating**: Establish standards for continuous model updating and adaptation while maintaining regulatory compliance, including approaches to version control and change management. These standards should enable AI systems to evolve in response to changing fraud patterns while ensuring that updates are properly validated and documented. Participants emphasized the importance of maintaining audit trails for model updates, enabling stakeholders to understand how and why AI systems change over time.
2. **Pattern Recognition Across Cases**: Develop standards to support linking seemingly isolated fraud cases to identify larger patterns while preserving privacy, enabling more effective detection of organized fraud schemes. These standards should enable AI systems to identify connections between different cases without compromising patient privacy or provider confidentiality. Participants noted that sophisticated fraud operations often involve multiple providers or claims, requiring systems that can detect patterns across seemingly unrelated cases.
3. **Machine-Readable Rules**: Standardize approaches to encoding healthcare rules in machine-readable formats, such as HL7's Clinical Quality Language. This

action would enable AI systems to check solutions against established standards, addressing the critical issue of AI hallucinations. These machine-readable rules would provide a foundation for validating AI-generated analyses, ensuring that they align with established healthcare contracts, policies, and regulations. Participants emphasized the importance of translating complex healthcare rules and gold standards of care into formats that can be used to validate both treatment decisions and payment claims.

Participants also discussed the importance of aligning AI implementation with existing regulatory frameworks, while also advocating for evolution of these frameworks to address emerging AI capabilities. This dual approach would enable healthcare organizations to deploy AI systems effectively within current regulatory constraints, including federal and state-level specific policies, while also working toward more supportive regulatory environments in the future.

HL7 Standards Support for Sensitive Data and Machine Rules

The HL7 Data Segmentation for Privacy (DS4P) standard and 'SecurityLabel' metadata value sets are being extended to represent AI-derived sensitivity tiers. These labels persist via FHIR 'Provenance', enabling downstream systems to enforce sharing policies.

For machine readable rules, Clinical Quality Language (CQL) and Clinical Decision Support (CDS) Hooks already support computable logic. Encoding CMS coverage rules or fraud detection heuristics as CQL libraries referenced from FHIR "PlanDefinition' resources lets AI services execute uniform rule sets across organizations.

5. Implementation Strategies and Pilot Programs

5.1 Provider-Payer Collaboration Pilots

HL7 Standards Support for Provider-Payer Pilots

Pilots involving providers and payers can employ SMART-on-FHIR apps and CDS Hooks services. These can be deployed in pilot sites to exercise new AI-traceability profiles. Bulk FHIR access can supply claims-review datasets. This could produce 'DetectedIssue' resources returned by an AI service reviewing and flagging suspect healthcare claims.

Participants strongly recommended developing pilot programs focused on provider-payer collaboration with streamlined medical record access and faster, more accurate upfront payment validation. These pilots would serve as proof-of-concept for the standards and approaches developed through HL7, demonstrating their practical value in real-world healthcare settings. Specifically, these pilots should:

1. **Structure a "win-win" scenario** where both provider and payer come to the table and collaborate. Participants emphasized that successful collaboration requires clear benefits for both parties, with providers gaining faster, more accurate payment and reduced administrative burden, and payers gaining better fraud detection, improved payment accuracy, and potentially lower administrative costs. One participant proposed a pilot focusing on collaboration between a hospital and a payer where medical records flow freely, enabling a lower-friction way to get providers paid more accurately upfront.

2. **Focus on specific high-volume claim types** where current processes create friction. By targeting areas with significant volume and known challenges, pilots could demonstrate measurable improvements in both payment accuracy and administrative efficiency. Participants suggested focusing on common procedures or diagnoses where billing errors are frequent or where current review processes are particularly cumbersome.

3. **Measure outcomes in terms of both payment accuracy and administrative efficiency**. Effective pilots would track multiple metrics, including reduction in payment errors, time to payment, administrative costs, and stakeholder satisfaction. These comprehensive measures would help demonstrate the full value of collaborative approaches, beyond just fraud detection.

Participants also discussed potential funding sources for these pilots, including venture capital firms that are actively investing in healthcare innovation. By engaging potential

funders early in the process, HL7's effort could help ensure that promising pilot concepts have the resources needed for successful implementation and scaling beyond the initial participants.

5.2 AI Validation Frameworks

HL7 Standards Support for AI Benchmark Validation

To support standardization of AI benchmark validation, HL7 could extend or create specific FHIR implementation guides or resources (such as structured 'MeasureReport' profiles or standardized library resources containing CQL logic to represent and share AI benchmark datasets. Such datasets would facilitate standardized validation across diverse healthcare environments, ensuring more consistent and trustworthy AI models.

Participants discussed the need to validate different AI methods applicable to healthcare from expert systems to machine learning.

The discussions yielded support for creating frameworks for implementing AI-assisted medical record review that adheres to both privacy standards and regulatory requirements:

1. **Checks and Balances Systems**: Develop systems that combine generative AI capabilities with rules-based verification mechanisms to ensure healthcare AI solutions are trustworthy and accurate. These systems would leverage the pattern recognition and natural language processing capabilities of generative AI while incorporating rules-based verification to prevent hallucinations or errors. Participants emphasized that effective AI systems for healthcare fraud detection would need to balance flexibility and creativity with reliability and consistency.
2. **Business Process Modeling Standards**: Create standards that translate complex healthcare rules and gold standards of care into formats that can be used to validate both treatment decisions and claims payment. These standards would enable more consistent application of healthcare policies, contracts, and regulations across different settings and systems. Participants noted that current approaches to encoding healthcare rules are often inconsistent or incomplete, creating challenges for both providers and payers.
3. **Validation Datasets**: Establish standardized datasets for validating AI models used in healthcare fraud detection, ensuring these datasets represent the diversity of healthcare delivery contexts. These validation datasets would enable more consistent evaluation of AI systems, facilitating comparison between different approaches and ensuring that systems perform effectively across diverse healthcare settings. Participants emphasized the importance of including edge

cases and challenging scenarios in these validation datasets, not just typical or straightforward cases.

Collectively these framework would provide a foundation for evaluating and comparing different AI approaches, helping healthcare organizations make informed decisions about which systems to implement.

5.3 Stakeholder Engagement and Education

Participants emphasized the importance of engaging key stakeholders in both the processes of standards development and implementation. This included different activities where private and public sector stakeholders could partner with HL7, including the following actions that could further advance standards-based approaches to AI in healthcare:

HL7 Standards Support for Clinical Definitions

HL7's CDS Hooks specification enables workflow-integrated interaction with Clinical Decision Support (including AI). Organizations can work with HL7 to encode granular clinical definitions, diagnostic criteria, and medical necessity standards directly into computable, shareable artifacts. Such artifacts would reduce variability in clinical interpretations and improve the accuracy of AI-driven claims validation.

1. **Working Groups**: Establish working groups with key stakeholders (providers, payers, and technology experts) to address standardization of medical record sharing and payment validation processes. These working groups would ensure that standards development is informed from multiple perspectives and is successful in addressing real-world challenges. Participants suggested including representatives from various healthcare settings, including rural and underserved areas that may face unique challenges.

2. **Policy Recommendations**: Explore policy recommendations to standardize clinical definitions and diagnostic criteria across the healthcare ecosystem. These recommendations would address underlying inconsistencies in how healthcare services are defined and categorized, which contribute to payment errors and disputes. Participants noted that current variations in clinical definitions create unnecessary complexity and confusion in healthcare billing and payment.

3. **Federal and State Alignment**: Engage with CMS and its program integrity teams to align initiatives with federal and state-level healthcare fraud prevention efforts. This engagement would help ensure that standards developed through the HL7 consensus process complement and support existing regulatory frameworks

and payment approaches rather than creating parallel or conflicting approaches. CMS has shown interest in innovative approaches to healthcare fraud prevention, creating opportunities for productive collaboration.

4. **Education and Training for Healthcare Stakeholders:** Provide education and training for healthcare stakeholders, ensuring that they understand how to effectively implement and use AI systems for fraud detection and prevention. This education would need to address both technical aspects of AI implementation and practical considerations for integrating these systems into existing workflows and processes.

5. **Cross-Sector Collaborative Pilots**: Establish collaborative pilot programs that bring together healthcare providers, payers, technology vendors, and regulatory bodies to test AI-enabled fraud detection and payment integrity solutions in real-world settings. As noted earlier in this report, the participants emphasized that successful pilots should include different healthcare settings (urban/rural, large/small organizations) and should measure outcomes in terms of both fraud reduction and administrative efficiency. Collectively these pilots would serve as "learning laboratories" where stakeholders can evaluate the effectiveness of different approaches, identify implementation challenges, and develop best practices for broader adoption. Participants emphasized that successful pilots should include diverse healthcare settings (urban/rural, large/small organizations) and should measure outcomes in terms of both fraud reduction and administrative efficiency. Results and lessons learned from these pilots should be widely shared through case studies, implementation guides, and peer-reviewed publications to accelerate adoption of effective approaches across the healthcare ecosystem.

6. Conclusion and Next Steps

The HL7 discussions on reducing fraud and improving payment integrity in healthcare through AI revealed both significant challenges and promising opportunities. The discussions highlighted that effective solutions must balance technical capabilities with practical implementation considerations and must create value for all stakeholders in the healthcare ecosystem.

6.1 Three Key Findings

Three key findings from these convenings and discussions include:

1. **AI has significant potential to transform healthcare fraud detection and payment integrity, while also improving the transition to value-based care**. By significantly reducing the cost of medical record review, AI could enable more comprehensive review of claims, leading to better detection of both unintentional errors and deliberate fraud. By improving the ability to tackle the complexity associated with payments associated with value-based care models, AI could advance the adoption of such models across the healthcare system.
2. **Successful implementation requires thoughtful standards development and frameworks that address explainability, transparency, data governance, and validation**. These standards must balance the need for innovation and flexibility with requirements for reliability, fairness, and regulatory compliance. Participants emphasized that standards development should lead to actionable outcomes rather than just theoretical frameworks.
3. **Pilot programs focused on provider-payer collaboration represent a promising path forward**. By demonstrating the practical value of collaborative approaches in real-world healthcare settings, these pilots could build momentum for greater administrative efficiency, more accurate billing from and reductions in unintentional errors by providers, and broader adoption of AI-enabled fraud detection and prevention systems among payers.

6.2 Recommendations for Immediate Action

1. **Commence working groups to develop standards or frameworks for explainable AI, transparency, data segmentation, and validation**. These

working groups should include diverse stakeholders, ensuring that standards development is informed by multiple perspectives and addresses real-world challenges.

2. **Initiate planning for pilot programs focused on provider-payer collaboration to improve payment integrity**. This planning should include identification of potential participants, development of specific use cases, and exploration of funding sources. Participants suggested engaging with potential funders early in the process to ensure that promising pilot concepts have the resources needed for successful implementation.
3. **Engage with federal and state agencies to ensure alignment with existing and emerging regulatory and payment approaches**. This engagement should include both federal agencies like CMS and state-level agencies, including Medicaid. Participants emphasized the importance of working within existing frameworks while also advocating for the evolution of these frameworks to better support AI implementation.

6.3 Envisioned Outcomes

The long-term recommended vision emerging from the discussions is a healthcare system where:

1. **Trusted AI tools are adopted to improve payment integrity and identify potential fraud while simultaneously streamlining valid claims processing**. These tools would enable more comprehensive review of claims at lower cost, leading to better detection of both unintentional errors and deliberate fraud. Participants emphasized that AI should be viewed as a tool to augment human expertise rather than replace it entirely, with clear frameworks for human-in-the-loop validation.
2. **Providers and payers collaborate effectively around payment, supported by interoperable systems and standards to reduce administrative burden and improve payment accuracy**. This collaboration would create benefits for all stakeholders in the healthcare ecosystem. Participants emphasized that effective collaboration requires alignment of incentives, with clear benefits for both providers and payers.
3. **AI models and related standards reflect and inform evolving regulatory requirements and payment models, including value-based care**. By developing standards that support more accurate and efficient gathering of data to support value-based payments, the healthcare system could align

incentives and better focus resources on interventions that deliver meaningful patient outcomes. Participants emphasized that this transition would require not just technical standards but also cultural and organizational changes.

The time for transformative action in healthcare payment integrity is now. By harnessing the power of AI through thoughtful standards development, collaborative pilot programs, and human-centered implementation strategies, we have an unprecedented opportunity to create a healthcare system that works better for everyone.

By pursuing the recommendations outlined in this report, HL7 and its partners can help realize this vision, reducing fraud, improving payment integrity, and ultimately improving the value and cost of healthcare for all stakeholders.

As we move forward together—providers, payers, technologists, and policymakers—we can embrace this moment of technological possibility with both optimism and responsibility. By focusing on "win-win" solutions that benefit all stakeholders, we can transform what has often been at times a challenged relationship between payers and providers into a collaborative partnership focused on value.

The recommendations in this report provide a practical roadmap for reducing costs while improving care quality, ensuring that our healthcare dollars are spent on what truly matters: patient outcomes.

7. Curated Listed of Report Call Outs

Page 14, HL7 Standards Support for Humans in the Loop:

> Both private and public sector organizations can employ existing HL7 FHIR resources like 'Task', 'AuditEvent' and 'DetectedIssue' which can be leveraged to define standardized workflows for human-in-the-loop validation.
>
> Extensions or profiles to these resources could more clearly specify roles, responsibilities, audit requirements, and decision accountability within clinical or administrative contexts.

Page 16, HL7 Standards Support for AI Results and Traceability:

> Organizations can employ existing HL7 FHIR resources such as 'Provenance', 'AuditEvent', 'DiagnosticReport', and 'Observation' to encapsulate AI outputs with full traceability.
>
> HL7 Implementation Guides can profile these resources, while CQL/FHIRPath expressions in the Clinical Reasoning module ('Measure'/'MeasureReport') can formalize validation rules.

Page 18, HL7 Standards Support for Explainable AI Artifacts:

> For explainability, a FHIR 'Evidence' resource, paired with the emerging 'Justification' element pattern, can serialize the reasoning trace for an AI decision. Linking the AI output 'Observation' to an 'EvidenceVariable' instance clarifies which data features contributed to the conclusion.
>
> For AI error metrics, existing HL7 FHIR 'MeasureReport' and 'Observation' can be leveraged to introduce standardized error reporting formats. HL7 could explicitly define extensions or value sets for reporting AI model performance metrics (accuracy, specificity, sensitivity, error rates) within these resources, ensuring consistent and transparent reporting across AI implementations in healthcare fraud detection.

Page 19, HL7 Standards Support for Sensitive Data and Machine Rules:

> The HL7 Data Segmentation for Privacy (DS4P) standard and 'SecurityLabel' metadata value sets are being extended to represent AI-derived sensitivity tiers. These labels persist via FHIR 'Provenance', enabling downstream systems to enforce sharing policies.

For machine readable rules, Clinical Quality Language (CQL) and Clinical Decision Support (CDS) Hooks already support computable logic. Encoding CMS coverage rules or fraud detection heuristics as CQL libraries referenced from FHIR "PlanDefinition' resources lets AI services execute uniform rule sets across organizations.

Page 20, HL7 Standards Support for Provider-Payer Pilots:

Pilots involving providers and payers can employ SMART-on-FHIR apps and CDS Hooks services. These can be deployed in pilot sites to exercise new AI-traceability profiles. Bulk FHIR access can supply claims-review datasets. This could produce 'DetectedIssue' resources returned by an AI service reviewing and flagging suspect healthcare claims.

Page 21, HL7 Standards Support for AI Benchmark Validation:

To support standardization of AI benchmark validation, HL7 could extend or create specific FHIR implementation guides or resources (such as structured 'MeasureReport' profiles or standardized library resources containing CQL logic to represent and share AI benchmark datasets. Such datasets would facilitate standardized validation across diverse healthcare environments, ensuring more consistent and trustworthy AI models.

Page 22, HL7 Standards Support for Clinical Definitions:

HL7's CDS Hooks specification enables workflow-integrated interaction with Clinical Decision Support (including AI). Organizations can work with HL7 to encode granular clinical definitions, diagnostic criteria, and medical necessity standards directly into computable, shareable artifacts. Such artifacts would reduce variability in clinical interpretations and improve the accuracy of AI-driven claims validation.

8. Acknowledgements

The authors acknowledge and extend a sincere thank you to the multiple people who provided their expert input, non-attributed thoughts, and feedback as to this report both through formal interviews and information conversations with the members of the Strategic Initiative on AI to Counter Fraud.

All conversations as part of this HL7 AI Strategic Initiative were provided expressly on a non-attributed basis to provide a safe space for candid discussions about the issues needed to be addressed and solved.

A **heartfelt thank you** to all who shared their insights, expertise, and wisdom to inform this report and the recommended actions.


HL7
International