SHIFT

The Modern Payment Integrity Buyer's Guide: Essentials for Speed, Accuracy & Transparency

www.shift-technology.com

Matching the momentum

The U.S. healthcare payer landscape is defined by constant movement. In 2025, the industry processed 420 CPT[1] and 252 ICD-10[2] code updates, alongside a steady stream of revisions to National and Local Coverage Determinations (NCD/LCD) from CMS. New forms of care delivery arise, or surges in care trends push coverage boundaries or drive questions for plans. This creates a significant operational challenge: the speed at which policy or care delivery evolves often outpaces the speed at which payment systems can be updated.

When a health plan relies on manual processes or static rule libraries to manage payment accuracy, the time required to translate new policy into production-ready logic creates a new window of exposure. During this lag, improper payments continue to flow, and the resulting pay-and-chase recovery efforts become a growing challenge on the other side.

We've heard it year over year in Payment Integrity, the goal is to shift Payment Integrity from a reactive, retrospective function into a proactive, transparent intelligence engine. This guide illustrates how systems fueled by agentic workflows, Generative AI and transparency are used to close the velocity gap, enhance prospective accuracy, and address the hidden administrative costs associated with black-box decision-making. When it comes to evaluating your current vendor or finding a new one, you'll know where to dig in.

1. "AMA Releases CPT 2025 Code Set." American Medical Association, 30 Aug. 2024, www.ama-assn.org/press-center/ama-press-releases/ama-releases-cpt-2025-code-set.
2. "ICD-10 Codes." Centers for Medicare & Medicaid Services, www.cms.gov/medicare/coding-billing/icd-10-codes. Accessed 24 May 2024.

Pacing with policy

Prepay accuracy—identifying errors before a claim is paid—is a high-impact, efficient way to manage medical loss ratios (MLR). However, the effectiveness of a proactive payment integrity solution can be limited by its update frequency.

Limitations of static rule libraries

Traditional payment integrity models rely on hard-coded rule libraries. Explained in a non-technical way, when a new policy comes out or is updated, a vendor must manually interpret the change, write new code, test it against historical claims, and finally deploy it. This is traditionally done through rigid deployment cycles, lasting 6 weeks or more.

This delay results in improper payment leakage or false positive alerts from concepts lagging behind policy changes. Because the system is operating on out of date logic, it fails to catch errors that align with current policy, guidelines or clinical standards. On top of that, these static rules can be too rigid to account for plan-specific nuances, leading to high false-positive or appeal rates that require manual intervention.

A typical deployment cycle of 6 weeks or more results in leakage from lagging behind policy changes.

Shifting toward dynamic edit integration & generation

A more efficient approach involves going beyond this static approach, with dynamic edit integration and generation. By integrating direct feeds from NCCI, I/OCE, and other CMS sources, Generative AI can ingest policy updates and suggest preventative logic changes in near real-time. Everest Group is reporting that GenAI enabled PI programs are reporting up to 42% decrease in administrative costs in running reviews & investigations at scale.

Accelerated updates & deployment: Logic updates are processed and ready for review shortly after a policy change is announced. This ensures the prospective gate remains aligned with the latest regulatory environment.

Reduced maintenance: Automated ingestion lessens the burden on internal IT and reduces burden (or overhead of) clinical teams, who otherwise spend significant time tracking and implementing these updates manually.

Granular customization: Plans can apply custom logic to specific contracts or provider groups, moving away from off-the-shelf edits that can lead to inaccuracies.

Applying a dynamic approach to FWA: FWA concepts and scenarios can adapt to the same policy, care delivery or provider behavioral trend changes to ensure evolving or new suspicious activity is captured.

Transparency is key: moving from “no” to “why” in Payment Integrity

Close behind a false positive, a leading source of friction in the payer-provider relationship (or vendor-payer relationship) is the lack of transparency in claims editing or FWA decisions. When a claim is flagged or denied without a clear explanation, it has a ripple effect of administrative actions that increase the cost of doing business.

“Because it said so” doesn’t cut it

You wouldn’t trust a traffic ticket without a clear violation explanation. Black box, or unexplainable decision making, forces providers into a cycle of headache-inducing appeals with payers. Without solid, defensible rationale for a denial, the provider’s natural response is to challenge the decision. This creates a grueling back and forth where both parties spend significant clinical and administrative resources appealing or defending claims.

When a claim is flagged or denied without a clear explanation, it has a ripple effect of administrative actions that increase the cost of doing business.

Research into the hidden costs of black-box systems indicates that these opaque decisions often negate the overall goal the PI program was designed to achieve - cost savings. When internal teams have to spend hours defending a vendor’s unexplainable flag, the operational ROI is diminished.

The mandate for explainable AI

The move toward explainable AI is a financial necessity and critical for maintaining healthy relationships with a plan’s provider network. A modern PI system should provide clear, policy-backed rationale for every alert.

Defensible decisions: By providing the specific NCCI guideline, contract clause or supporting information that triggered an alert, the conversation with the provider moves from a dispute to a data-driven verification.

Streamlined provider communication: Transparency allows for quicker resolution of issues and a measurable reduction in appeal volumes - and even educational opportunities. When the rationale is clear, the provider is more likely to accept the decision or correct the billing behavior, reducing long-term friction.

Scaling speed & human expertise with agents

Agents are everywhere these days. While the hype cycle is in full swing, the operational reality is catching up: according to a recent McKinsey study, 62% of organizations are already experimenting with AI agents, and 23% have successfully started scaling them in at least one business function.[3] They have moved beyond a buzzword; agents are a clear process accelerator for both speed and complex analysis. By automating repetitive research and data-gathering tasks, agentic workflows allow clinical and investigative teams to bypass manual bottlenecks, ensuring that payment integrity updates and FWA detection keep pace with the shifting regulatory landscape.

While automation is essential for high-volume claims editing, certain complexities of claims editing as well as FWA detection require a human-in-the-loop approach for validation and optimization. The goal is to use AI to augment or accelerate team capacity, not to replace their judgment.

Defining agentic workflows

In the context of Payment Integrity, Agentic AI refers to autonomous decision support and identifying, suggesting or automating the next best action in a workflow. In Payment Integrity specialized agents designed to handle the heavy lifting of tasks such as data analysis, content summarization and collection, policy research, logic generation and more. See some sample areas agents can help in the payment integrity process.

Where agents can be created for any use case or repeatable task, here are a few examples of process accelerators within the Payment Integrity workflow:

Agent-led edit development

Rather than a reviewer manually searching through multiple databases and policy documents to generate an edit to validate a claim, an AI Agent performs the initial research. It identifies the relevant policy sections and surfaces the necessary data points for new logic to be created and deployed in a system.

Summarizing & packaging leads with an agent

Agent delivers an actionable Lead Package to the investigator or clinical coder. This package typically includes:

Data summarization: A consolidated view of the claim or provider analysis; provider's billing patterns and outlier or anomalous activity.

Contextual evidence: Direct links to the specific clinical guidelines or plan policies that suggest an error or behavioral shift.

Anomaly detection: Highlighting specific claim lines or other activity that deviate from expected patterns, allowing the reviewer to focus on high-value tasks.

This approach significantly accelerates claims editing and FWA workflows. By removing manual research and summarization steps, teams can manage a higher volume of cases with greater precision.

[3] The State of AI in 2025: Agents and Innovation. McKinsey & Company, Nov. 2024, www.mckinsey.com/~/media/mckinsey/business%20functions/quantumblack/our%20insights/the%20state%20of%20ai/november%202025/the-state-of-ai-2025-agents-innovation_cmyk-v1.pdf.

Speed-to-value: Streamlining implementation

A major first success of a PI initiative is often determined by the ease of implementation. Legacy systems that require extensive data mapping and rigid schemas create a pull on the plan's IT resources.

Limitless data integration

Embracing the "good data in, good data out"mantra, a modern payment integrity solution must deliver limitless data integration at scale, capable of ingesting diverse data sources without restrictions to format. This capability is critical, as it gives health plans and their PI systems the best view of claim activity while avoiding the extensive, high-resource demands of manual data cleansing and mapping that frequently stall implementation plans. A few things to consider below:

Native format ingestion: The ability to process claims, membership, and contract data without requiring a total overhaul of the plan's data architecture.

Automated mapping: Utilizing healthcare-trained AI to consistently label and align data, ensuring that the analytical results are accurate from the start.

Intelligent data cleansing: AI models to enrich data sources, resolving entities such as multiple records to increase accuracy and efficiency of analysis.

Single implementation for multiple deployments

Consolidation, or moving away from a stack of solutions, is a growing trend, but deciding how and when to consolidate can be daunting - especially after a long initial implementation process. When selecting a vendor, prioritize one whose implementation strategy supports future growth and solution scaling from a single source. This ensures that when the time comes to consolidate payment integrity needs you won't have to undergo multiple, separate implementations for different solution deployments.

The goal is to produce value as quickly as possible. By reducing the IT lift and utilizing automated implementation steps, a health plan can begin identifying savings within days after go-live. This efficiency allows the organization to realize an ROI much sooner than traditional implementation timelines allow.

Vetting for value: 5 questions for evaluation

When evaluating a potential PI partner, it is important to look beyond high-level feature lists and ask questions that address operational reality.

Update/production speed: "What is the specific lag time between a CMS update and that logic being live in production? Is this update performed manually or through AI-assisted ingestion?"

Transparency & explainability: "Can you demonstrate how your system surfaces the clinical or policy rationale for an edit or alert? Is this information available to our provider relations team in a clear, non-technical format?"

Traceability: "Is there a clear audit trail for every AI-driven decision? Can our internal auditors see the specific data points and policies used to generate a flag?"

Proactive detection: "How does your system identify and alert on behavior outside of a rules or scenario library? What mechanisms are in place to deploy new logic for novel detection? How do you suppress false positives as provider behavior evolves?"

Ease of implementation: "Are there data format or mapping requirements for data transfers and ingestion? What responsibilities or expectations does the health plan have in the implementation process? What implementation steps are required for other deployment types (pass position, FWA capabilities, etc.)"

Payment Integrity as a strategic asset

Payment Integrity has evolved from a back-office compliance function into a central component of a health plan's financial strategy. In an environment of increasing clinical complexity and tightening margins, the ability to ensure accurate payments in with ongoing alignment with policy and transparent outcomes is a significant competitive advantage.

By focusing on keeping edits and concepts current, explainable AI, and agentic decision support, health plans can protect their MLR while simultaneously reducing the administrative friction that drives up operational costs and strains provider relationships. Transitioning from a static, rule-based approach to a dynamic intelligence engine allows a plan to remain accurate, defensible, and responsive to the ever-changing healthcare landscape.

The focus must remain on speed, precision and transparency. In doing so, payment integrity becomes not just a tool for recovery, but a strategic asset for long-term operational health.

Visit us for more information on Shift Technology and Payment Integrity.