



The MLR Imperative

\$6B Cost Containment through Intelligent Medical Expense Management

Unlocking Hidden Efficiency in Payment Integrity
and Administrative Processes



Ashish Jaiman
Founder & CEO

November 2025



Executive Summary

Healthcare payers face an unprecedented financial squeeze. The Affordable Care Act's Medical Loss Ratio (MLR) requirements mandate that 80-85% of premium revenue be spent on medical care and quality improvement, leaving only 15-20% for administrative functions such as claims processing, payment integrity, member services, technology infrastructure, and profit.

With medical costs rising 6-8% annually and regulatory pressure intensifying, payers must extract maximum efficiency from every administrative dollar.

Yet hidden within this constraint lies a \$6B+ efficiency opportunity. Current payment integrity approaches consume 2-4% of total premium dollars while simultaneously creating provider friction that adds another 1-2% in hidden administrative costs (rework, appeals, disputes, provider inquiries). Meanwhile, payment leakage, the difference between what should be paid per policy (the allowed amount) and what is actually paid, exceeds \$100B annually across the industry, representing 3-5% of total medical expenses.

This white paper demonstrates how an AI-native payment integrity platform can simultaneously reduce both administrative burden and payment leakage, creating a dual pathway to MLR improvement.

By deploying Nēdl Labs's neuro-symbolic AI that builds provider trust rather than eroding it, payers can achieve 30-50% reductions in payment integrity operational costs while improving medical expense accuracy by 1-2 percentage points, translating to **\$6B+** in annual industry savings.

Key Findings

- **MLR squeeze creates urgency:** Payers spending >85% on medical costs are only left with <15% for all administrative functions. Every 1% improvement in administrative efficiency = \$2B+ industry opportunity
- **Hidden costs exceed direct costs:** Provider friction from aggressive payment integrity creates 1-2% of premium in hidden administrative burden through appeals, disputes, and rework
- **Dual efficiency pathway:** AI-native platforms reduce both administrative costs (30-50% operational savings) and payment leakage (1-2% MedEx improvement)
- **Trust builds efficiency:** Transparent AI that guides providers pre-submission reduces post-payment friction, converting administrative burden into preventive alignment
- **Regulatory alignment drives ROI:** CMS compliance requirements favor explainable AI, creating a competitive advantage for early movers while reducing litigation risk

EXECUTIVE SUMMARY	1
KEY FINDINGS	1
I. THE MLR SQUEEZE: UNDERSTANDING THE \$6B+ OPPORTUNITY	3
THE MATH OF MLR PRESSURE	3
HIDDEN ADMINISTRATIVE COSTS: THE PROVIDER FRICTION TAX.....	4
II. THE DUAL EFFICIENCY OPPORTUNITY: ADMINISTRATIVE + MEDICAL EXPENSE	5
PATHWAY 1: ADMINISTRATIVE EFFICIENCY THROUGH TRUST-BUILDING	5
PATHWAY 2: MEDICAL EXPENSE ACCURACY THROUGH BETTER DETECTION	6
PROVIDER AI EVOLUTION: OPPORTUNITY AND EXPLOITATION.....	7
III. NĒDL'S AI-NATIVE NEURO-SYMBOLIC PLATFORM: TECHNICAL ARCHITECTURE FOR DUAL EFFICIENCY.....	7
THREE-LAYER ARCHITECTURE	7
WHY NEURO-SYMBOLIC AI ENABLES DUAL EFFICIENCY.....	8
IV. FINANCIAL MODEL: QUANTIFYING THE \$6B+ INDUSTRY OPPORTUNITY.....	9
BASELINE ASSUMPTIONS (CURRENT STATE)	9
AI-NATIVE IMPACT (NĒDL PLATFORM IMPLEMENTATION)	9
MLR IMPACT	10
V. CONCLUSION: FROM CONSTRAINT TO COMPETITIVE ADVANTAGE.....	10
THE STRATEGIC IMPERATIVE.....	10
THE PATH FORWARD	11
VI. ABOUT THE AUTHORS.....	11
NĒDL LABS	11

The MLR Squeeze: Understanding the \$6B+ Opportunity

The Affordable Care Act's Medical Loss Ratio requires fundamentally reshaped healthcare payer economics. By mandating that 80-85% of premium revenue be spent on medical care and quality improvement activities, the ACA created a hard ceiling on administrative spending.

For a payer with \$50B in annual premium revenue, this means just \$7.5-10B is available for all non-medical costs: claims processing, member services, network management, payment integrity, technology infrastructure, broker commissions, regulatory compliance, and profit.

The Math of MLR Pressure

Consider the financial dynamics facing a typical large payer:

Revenue & Expense Category	Amount (Annual)
Total Premium Revenue	\$50,000,000,000
Medical Expenses (MLR: 85%)	\$42,500,000,000
Available for Administration (15%)	\$7,500,000,000
Claims Processing & Adjudication	\$2,000,000,000
Payment Integrity Programs	\$1,500,000,000
Member Services & Call Centers	\$1,000,000,000
Network Management & Provider Relations	\$800,000,000
Technology, Compliance, Sales, Profit	\$2,200,000,000

The challenge is stark: Payment integrity programs consume 3% of total premium (\$1.5B in this example) but often create hidden costs that exceed their direct expenses.

When aggressive denial systems force providers to appeal 15-20% of denials, and overturn rates reach 10-18%, the administrative burden extends to both payer appeals teams and provider revenue cycle staff.

\$100B+

Annual Payment Leakage (Commercial Plans)

Despite aggressive payment integrity programs and increasing AI deployment

Hidden Administrative Costs: The Provider Friction Tax

Traditional cost accounting captures direct payment integrity expenses, such as software licenses, staff salaries, and vendor fees. But it misses the cascading costs created by adversarial payment integrity approaches:

Direct Administrative Costs (Visible)

- Payment integrity software and vendor fees: \$300-500M annually
- Staff costs for pre-payment review, post-payment audit, appeals processing: \$600-800M
- Claims rework from denial-resubmission cycles: \$200-300M

Total visible costs: \$1.1-1.6B (2.2-3.2% of premium)

Hidden Administrative Costs (Invisible)

- **Provider inquiry volume:** When denials lack clear explanations, providers call payer customer service. Cost: \$150-250M annually in additional call center burden
- **Network contracting impact:** Aggressive payment integrity becomes a negotiating point in provider contracts. High-value specialists demand higher rates to offset the perceived risk of underpayment. Cost: \$200-400M in rate pressure
- **Provider churn and replacement costs:** Dissatisfied providers leave networks, requiring expensive recruitment. Cost: \$100-200M
- **Member satisfaction impact:** Billing disputes and coverage confusion from opaque denials reduce CAHPS scores, affecting Star Ratings and quality bonuses. Cost: \$50-150M in lost bonus payments
- **Litigation and regulatory scrutiny:** Class action lawsuits and regulatory investigations. Cost: \$100-200M in legal defense and settlements

Total hidden costs: \$600-1,200M (1.2-2.4% of premium)

Combined Reality: For every dollar spent directly on payment integrity, payers incur \$0.50-\$0.75 in hidden administrative costs created by provider friction.

A \$1.5B payment integrity program actually consumes \$2.25-2.6B in total administrative capacity, representing 4.5-5.2% of premium versus the 3% visible in budgets.

The Dual Efficiency Opportunity: Administrative + Medical Expense

Most payers approach MLR optimization through a single lens: reduce administrative costs. But this creates a false trade-off between cost efficiency and payment integrity effectiveness. The breakthrough insight is that **the same AI-native technologies that reduce administrative burden also improve medical expense accuracy**, creating a dual pathway to MLR improvement.

Pathway 1: Administrative Efficiency Through Trust-Building

Traditional payment integrity operates post-payment: claims are submitted, algorithms flag issues, denials are issued, providers appeal, payers review appeals. Each step consumes administrative resources from both parties. The cycle repeats millions of times annually.

AI-native platforms shift this dynamic from post-payment friction to pre-submission alignment:

Traditional Approach (Post-Payment)	AI-Native Approach (Pre-Submission)
Provider submits claim	Provider validates claim pre-submission
↓	↓
Payer algorithm flags an issue	AI provides real-time feedback.
↓	↓
Denial issued (often vague)	The system specifies exact requirements.
↓	↓
Provider calls to understand	Provider supplements documentation
↓	↓

Traditional Approach (Post-Payment)	AI-Native Approach (Pre-Submission)
Provider appeals with evidence. ↓ Payer reviews the appeal ↓ 10-18% overturn rate	Claim approved on first submission. ↓ No appeals cycle needed ↓ 72% reduction in appeals
Administrative Burden: 7-10 touches per claim 45-90 day payment cycles Provider frustration & inquiries	Administrative Efficiency: 2-3 touches per claim 15-30 day payment cycles Provider trust & alignment

Administrative savings breakdown:

- Appeals processing staff reduction: 40-50% (7-10 fewer appeals per 100 claims)
- Call center volume reduction: 30-40% (fewer provider inquiries about denials)
- Claims rework reduction: 50-60% (clean claims on first submission)
- Payment integrity technology costs: 20-30% reduction (automated real-time guidance vs. batch post-payment audit)

30-50% reduction in total payment integrity operational costs, freeing \$450-750M annually for a \$50B premium payer.

Pathway 2: Medical Expense Accuracy Through Better Detection

While administrative efficiency improves the denominator in MLR calculations, medical expense accuracy improves the numerator. Current payment integrity approaches achieve 2-5% audit sampling rates, missing systematic patterns that AI-native platforms detect through comprehensive behavioral monitoring.

The opportunity: If payment leakage represents 3-5% of medical expenses (\$1.3-2.1B for a \$50B payer with \$42.5B in MedEx), improving detection by just 20-40% captures \$260-840M annually. Combined with administrative savings, this creates the dual efficiency pathway.

Provider AI Evolution: Opportunity and Exploitation

Understanding the provider AI landscape is essential for designing effective defensive strategies. Provider adoption of AI tools for claims optimization is accelerating rapidly, driven by economic necessity (denial rates up 37%), labor shortages (70% of organizations facing staffing challenges per Experian data), and technology maturation (AI-powered revenue cycle management now mainstream).

nēdl's AI-Native Neuro-Symbolic Platform: Technical Architecture for Dual Efficiency

Achieving both administrative efficiency and medical expense accuracy requires a fundamentally different technical architecture than traditional payment integrity systems. nēdl Pulse implements neuro-symbolic AI specifically designed for the dual efficiency pathway.

Three-Layer Architecture

Layer 1: Policy Intelligence Infrastructure (Eliminates Administrative Friction)

Problem: Provider friction stems from opacity. When providers don't understand why claims are denied, they must reverse-engineer rules through trial and error, creating a massive administrative burden for both parties.

Solution: Convert all policies (NCDs, LCDs, medical policies, contracts) into machine-readable knowledge graphs. Every adjudication decision traces to specific policy language with complete explainability.

Impact on MLR:

- Providers validate claims pre-submission against actual policies → 50-70% reduction in unclear denials
- Provider inquiry calls drop 35-45% → call center cost savings of \$50-110M
- Appeals overturn rate drops from 10-18% to 3-5% → appeals processing cost savings of \$200-350M

Layer 2: Behavioral Pattern Detection (Improves MedEx Accuracy)

Problem: Traditional audit samples 2-5% of claims. Bad actors exploit low sampling rates by distributing questionable claims to avoid detection triggers.

Solution: Continuous monitoring of all providers and claims using neuro-symbolic AI that combines statistical pattern detection with clinical reasoning. AI flags claims where documentation meets technical requirements but doesn't align with standard treatment protocols.

Impact on MLR:

- Detection rates increase 5-10x vs. traditional sampling → identifies \$260-840M in additional payment leakage
- False positive rate drops 60-75% → legitimate providers face less scrutiny, reducing friction
- Clinical incongruence scoring focuses audits on the highest-risk claims → audit staff productivity increases 40-60%

Layer 3: Real-Time Collaborative Adjudication (Optimizes Both)

Problem: Post-payment adjudication creates an adversarial dynamic where providers optimize for approval rather than accuracy, and payers optimize for denial rather than appropriate payment.

Solution: Provider portal with pre-submission validation. The system provides immediate feedback on policy compliance, specifies evidence requirements, and suggests alternatives when claims won't be approved.

Impact on MLR:

- First-pass approval rates increase 15-25% → reduces claims rework costs by \$150-225M
- Payment cycle time decreases 40-60% → improves provider satisfaction and cash flow
- Provider optimization aligns with clinical appropriateness → MedEx accuracy improves without administrative burden

Why Neuro-Symbolic AI Enables Dual Efficiency

Traditional machine learning cannot achieve this dual pathway because it lacks explainability. Providers cannot align with opaque algorithms, and regulators cannot verify that AI decisions reflect individual patient assessment. Neuro-symbolic architecture solves both problems:

- **For administrative efficiency:** Explainable decisions enable provider self-service validation, eliminating post-payment friction
- **For MedEx accuracy:** Clinical reasoning identifies incongruence that pure statistical models miss, improving detection precision
- **For CMS compliance:** Individual patient assessment with transparent reasoning satisfies regulatory requirements

Financial Model: Quantifying the \$6B+ Industry Opportunity

The following model demonstrates how AI-native payment integrity creates dual efficiency gains for a typical large payer with \$50B in annual premium revenue. Results scale proportionally for payers of different sizes.

Baseline Assumptions (Current State)

Metric	Annual Amount
Total Premium Revenue	\$50,000M
Medical Expenses (85% MLR)	\$42,500M
Payment Integrity Direct Costs (3% of premium)	\$1,500M
Hidden Administrative Costs from Provider Friction (1.5% of premium)	\$750M
Payment Leakage (4% of MedEx)	\$1,700M

AI-Native Impact (nēdl Platform Implementation)

Improvement Area	% Improvement	Annual Savings
Payment Integrity Direct Cost Reduction	40%	\$600M
Hidden Administrative Cost Elimination	65%	\$488M
Payment Leakage Reduction	30%	\$510M
Total Annual Financial Impact	—	\$1,598M

Industry Extrapolation: \$6B+ Annual Opportunity

At \$1.6B in savings per \$50B payer, the addressable US market (roughly \$1.2T in commercial + MA premium) represents \$38B+ in total inefficiency, of which 15-20% is immediately capturable through AI-native platforms, creating a \$6-8B annual industry opportunity.

MLR Impact

- **Administrative Efficiency:** \$1.088B in admin cost savings = 2.2% of premium freed for profit or invested in member benefits
- **Medical Expense Accuracy:** \$510M leakage prevention = 1.2% improvement in MedEx efficiency, directly improving MLR
- **Combined Effect:** 3.4 percentage point improvement in effective MLR management, creating sustainable competitive advantage in a tight margin environment

Conclusion: From Constraint to Competitive Advantage

The MLR squeeze, once viewed purely as a regulatory constraint, has become a catalyst for a fundamental transformation in payment integrity. Payers who continue deploying traditional approaches will face escalating administrative costs, deteriorating provider relationships, and regulatory scrutiny. Those who embrace AI-native platforms will unlock a \$6B+ efficiency opportunity by simultaneously improving administrative efficiency and medical expense accuracy.

The Strategic Imperative

Three converging forces make this transformation urgent:

- **Economic pressure:** Medical costs are rising 6-8% annually while MLR requirements remain fixed. Every year of delay increases the administrative burden relative to the available budget.
- **Provider AI adoption:** As providers deploy their own AI optimization tools, payers without sophisticated defensive AI face accelerating payment leakage.
- **Regulatory evolution:** CMS requirements for explainable AI will only tighten. Early movers establish a competitive advantage before mandates force expensive remediation.

The Path Forward

Realizing the dual efficiency opportunity requires both technological capability (nēdl's neuro-symbolic AI platform) and implementation expertise

Recommended Approach:

1. **90-day pilot:** Deploy in high-volume, high-denial category. Measure administrative cost reduction and MedEx improvement. Target: 30-40% denial reduction, 50-60% appeals reduction.
2. **180-day expansion:** Scale to 3-5 additional categories. Implement behavioral detection. Deploy the provider portal. Target: Demonstrate \$200-400M annual run-rate impact.
3. **12-month full deployment:** Complete policy digitization across all categories. Full integration with existing systems. Target: Achieve \$1.2-1.8B annual impact, 12-18 month ROI payback.

The MLR imperative is not a constraint to manage—it's an opportunity to lead.

Payers who deploy trust-first defensive AI will simultaneously reduce costs, improve accuracy, enhance provider relationships, and satisfy regulators. Those who delay will face escalating costs defending an unsustainable status quo.

*The \$6B+ efficiency opportunity is real. Technology is built. The question is:
Who will capture it?*

About the Author

nēdl Labs

nēdl Labs is pioneering AI-native payment integrity solutions for healthcare payers. Our neuro-symbolic AI platform combines neural networks' pattern recognition with symbolic reasoning's explainability, enabling payment integrity systems that simultaneously reduce leakage and build provider trust.

nēdl brings deep expertise in responsible AI, healthcare policy, and enterprise product development to the payment integrity challenge.

Contact: contact@nedllabs.com

Web: nedllabs.com



© **2025 nēdl Labs. All rights reserved.**

This white paper is provided for informational purposes only. No part of this publication may be reproduced without written permission.

The information contained herein is subject to change without notice. This document is not a commitment to provide any particular product, feature, or service.