innovaccer

HAP E-BOOK SERIES 3

The ROI of Autonomous Healthcare Operations

Unlocking Efficiency, Scale, and Better Outcomes

An Innovaccer Gravity Executive Brief for CEOs, CFOs, and COOs of US Health Systems

The Business Case Most CFOs Haven't Seen Yet

Healthcare CFOs are being asked to reduce administrative costs while managing rising labour costs, growing denial rates, an expanding regulatory burden, and a workforce gap that cannot be closed by hiring alone.

The standard response has been AI investment. The standard result has been disappointment.

Not because the models are bad. Because 74% of healthcare AI is operating on fragmented data foundations, and the ROI never materialises when agents are confident and wrong.

This brief makes the financial case for a different approach, and introduces a commercial model that removes the investment risk entirely.

Why the Current ROI Model Has Stopped Working

Most healthcare AI investments follow the same logic: invest in a point solution, measure productivity improvement per FTE, calculate savings against labour cost.

This model has a structural ceiling.

WHY FRAGMENTED AGENTS DESTROY ROI

The efficiency gain from faster processing is consumed by the rework generated by incorrect processing

01	02	03	04
Fragmented agent acts	**Incorrect output**	**Human rework required**	**Efficiency gain consumed**
No clinical history. No payer context. Operates on a slice.	Wrong approval. Wrong code. Wrong claim. Delivered confidently.	$25 to $35 per claim. Review, correct, resubmit.	Speed savings erased. Net ROI approaches zero.

The fix is not faster agents. It is accurate agents, operating on a unified data foundation that eliminates the rework at the source.

When agents operate on fragmented data, they make mistakes that create rework. A prior auth agent without clinical history approves the wrong thing, generating a downstream denial that costs $25 to $35 to rework. A coding agent without the signed note codes incorrectly, requiring human review, correction, and resubmission.

The efficiency gain from faster processing is consumed by the rework generated by incorrect processing.

The right ROI model is not faster agents. It is accurate agents, operating on a unified clinical, financial, and operational data foundation that eliminates the rework at the source.

The Commercial Model That Changes the Equation

Before discussing ROI projections, it is worth understanding how HAP's commercial model changes the risk calculus.

Standard enterprise software invoices regardless of performance. Whether the platform delivers or not. Whether staff use it or not. Whether the outcomes promised materialised or not. The vendor's revenue is locked in at signing. After that, your outcomes are your problem.

HAP is priced on outcomes: prior authorisations processed, denials overturned, charts reviewed accurately. If Innovaccer underperforms, it does not get paid.

TWO VERY DIFFERENT COMMERCIAL MODELS

How HAP's outcome-based pricing removes the investment risk entirely

Standard enterprise model	Outcome-based model (HAP)
WHEN THE INVOICE ARRIVES At contract signing, regardless of deployment progress or outcomes	WHEN THE INVOICE ARRIVES When outcomes are delivered: auths processed, denials overturned, charts reviewed
WHO BEARS PERFORMANCE RISK The health system. Vendor revenue is locked in at signing.	WHO BEARS PERFORMANCE RISK Innovaccer. If the platform underperforms, it does not get paid.
IF OUTCOMES MISS Invoice still arrives. Contract still enforced. Your problem to solve.	IF OUTCOMES MISS No payment. Both parties remain on the same side of the table.
VENDOR INCENTIVE AFTER SIGNING Renewal. Not your performance.	VENDOR INCENTIVE AFTER SIGNING Your outcomes. Every month.

For a CFO evaluating AI investment, this changes the risk model entirely. The downside is removed. The question shifts from "will this deliver?" to "how fast do we expand?"

Three Dimensions of ROI

Why foundation-first autonomous administration compounds over time

DIMENSION 1

Direct cost reduction

Measurable from Day 1

Mechanism

Accurate agents eliminate rework at the source. Prior auth first-pass rates move from 70 to 75% up to 85 to 92%. Denials prevented before claims go out rather than recovered after. Risk adjustment captured continuously, not periodically.

$840K+

avoided rework per year from prior auth alone at 150K auths

DIMENSION 2

Operating leverage

Compounds with scale

Mechanism

Volume growth does not require proportional headcount growth. Attrition in RCM stops driving capacity loss. New use cases deploy on the same foundation rather than as new integration projects.

25–35%

annual RCM attrition no longer driving capacity loss

DIMENSION 3

Capacity dividend

Rarely modelled, most significant

Mechanism

Nurses freed from 35% administrative burden return that time to patients. RCM specialists move from processing to judgment. Skilled staff most expensive to replace choose to stay.

35%

of nursing shift returned to patient care

Dimension 1: Direct cost reduction

The fully loaded cost of a prior authorisation runs $30 to $80, covering staff time, rework, peer-to-peer reviews, and tracking. When HAP's prior auth agent operates on unified clinical and payer data rather than just the request, it acts correctly from first submission. First-pass approval rates move from the industry average of 70 to 75% up to 85 to 92%.

For a health system processing 150,000 prior auths annually, that means 24,000 additional first-pass approvals and $600K to $840K in avoided rework per year, from prior auth alone.

US health systems spend approximately $260 billion annually processing denied claims, at an average rework cost of $25 to $35 per claim. HAP's denial management agent does not wait to be triggered. It knows a denial is likely before the claim goes out, because it can see the full clinical and payer picture. Prevention is economically superior to recovery in every dimension.

Dimension 2: Operating leverage

When agents operate on a unified foundation, volume growth does not require proportional headcount growth. The attrition problem in RCM, running at 25 to 35% annually, stops driving capacity loss. New use cases deploy on the same foundation rather than as new integration projects.

A health system growing 15% in net revenue over five years while holding administrative cost relatively flat through autonomous operations demonstrates the operating leverage that bond raters, investors, and board members recognise.

Dimension 3: The capacity dividend

This dimension is rarely modelled but is often the most significant.

When nurses are freed from 35% administrative burden, that time goes to patients. When RCM specialists stop processing routine prior auths, their capacity goes to complex case management, payer relationships, and decisions that genuinely require their expertise.

That capacity dividend does not appear in a traditional ROI model, but it drives patient outcomes, clinician satisfaction, and the retention of the skilled staff most expensive to replace.

Building the Business Case

1. **Step 1: Establish the true baseline.**

   Most health systems find their true administrative cost is 20 to 30% higher than the operating budget when they include full labour cost with turnover and retraining, technology cost across all point solution licences and integrations, revenue leakage from unappealed denials and unrecovered underpayments, and the rework generated by agents operating on incomplete data.

2. **Step 2: Model on accurate agents, not just faster ones.**

   The ROI of foundation-first autonomous administration comes from accuracy, not speed. Model the impact of moving first-pass approval rates from 72% to 88%. Model the upstream prevention of denials rather than downstream recovery.

3. **Step 3: Model the five-year compounding benefit.**

   Every case HAP processes makes the platform smarter. Year 3 ROI is higher than Year 1, not because the platform deploys faster, but because it has processed more cases, learned more patterns, and built deeper payer intelligence.

Documented Outcomes at Scale

These are not projections. They are outcomes delivered at health systems currently live on the platform.

$7M to $19M in annual savings from technology consolidation and administrative cost reduction. 23% improvement in prior auth first-pass approval rates. 8% increase in PCP visit volume through improved patient access workflows. 7% increase in revenue capture through autonomous risk adjustment. 40 days average from use case conception to deployment. $1B+ in total documented administrative savings delivered.

See the ROI Before You Commit

The AI Bootcamp eliminates the evaluation risk. Innovaccer's technical team comes onsite, learns your specific workflows, and builds working agents against your real data in 48 hours. You see real outcomes, including approval rates, processing times, and revenue recovered, from your actual environment.

You see the ROI before you commit. The outcome-based pricing ensures you only pay when it continues.

REQUEST YOUR AI BOOTCAMP AT INNOVACCER.COM/HAP

About Gravity by Innovaccer

Gravity is the Healthcare Autonomy Platform, built foundation-first.

14 years of unified clinical, financial, and operational intelligence. 200+ EHR connectors. 6,000+ data quality rules. 80M patient lives. Every agent inherits that foundation. Every agent's learning enriches every other.
Priced on outcomes: Innovaccer earns when you do.

Trusted by 7 of the top 10 US health systems · $1B+ in documented savings · 8 weeks to production · Best in KLAS 2026