ANALYST PERSPECTIVE

Continuous Audit as a Strategic Asset

Lessons from Cybersecurity for Healthcare Procurement

WHY THIS PAPER, WHY NOW

Ask a CFO how much their health system spends with its top twenty suppliers and you will likely get three different numbers — one from AP, one from the ERP, one from the contracts team. All three can be challenged. The vendor master that sits beneath every one of those answers has never been cleaned to the standard any other risk-bearing discipline would accept. Cybersecurity faced the same problem a decade ago and solved it by building Continuous Controls Monitoring as a named category. Procurement has not made that move. This paper makes the case that it must — and shows what the operating model looks like when it does.

AUDIENCE

CFO · CPO · Chief Audit Executive · Chief Data Officer · Healthcare provider boards

AUTHORS

Dr. Ramyaa Ganesh, Dr. Vegi S.

CO-AUTHORS

Jeffrey Heenan-Jalil, Svetlana Schenkova, Hunter Heenan-Jalil, Valeria Artsybasheva

PUBLISHED

June 2026 · hunterAI Analyst Perspective Series · Confidential

ANALYST PERSPECTIVE

FOREWORD by John Andrews

Over the past two decades, I have witnessed healthcare supply chains evolve from transactional purchasing and back-office operations to strategic functions that can support and enhance a health system's financial and clinical performance. Despite these advancements, one thing remains constant: our decisions are only as good as the data beneath them.

This white paper names something many of us spent years working around: the constraint was rarely the tools we bought and implemented, but the integrity of the data underneath them. That data, the vendor master beneath every PO, every contract, and every payment, is now a potential risk in the modern health system. We have invested heavily in spend analytics, AI, contract lifecycle management, strategic sourcing, and purchased-services optimization, but rarely with parity in the foundation that drives and supports all of it.

What remains is a paradox we can all recognize. We negotiate effectively but report spend inconsistently. We leverage advanced analytics but trust the outputs selectively. We are often reactive instead of proactive, addressing rogue spending and spend leakage after it happens rather than preventing it in the first place.

Cybersecurity faced a similar dilemma a decade ago. It responded by naming the control, Continuous Controls Monitoring, and by treating data integrity as a standing discipline rather than a project. As the paper points out, procurement and AP have yet to fully make that journey.

What this paper makes clear is that the economic realities in healthcare have shifted. With margins compressed and AI more likely to amplify defects in the master data than to catch them, governing this data continuously can no longer be a "nice to have." It should be a part of the foundation that responsible financial stewardship and supply chain excellence are built on.

For boards, CFOs, and supply chain leaders, the message is straightforward. A 21st-century health system cannot run on yesterday's data governance. Still, the same discipline the paper brings to the problem belongs in the solution: weigh its economics against your own data before acting on them.

The central premise is straightforward: vendor and supplier data must be governed as a continuous, active discipline, reviewed and audited on an ongoing basis, rather than handled as a periodic project. The research the authors present points to a real cost in neglecting this discipline, and that is exactly why their argument lands. Treating it as standard practice is, to my mind, simply part of responsible financial stewardship and supply chain excellence.

John Andrews
Adjunct Faculty Health Systems Management College of Health Sciences Rush University

Contents

Executive Summary

More than 30 cents of every operating dollar a healthcare provider spends goes through vendor master, accounts payable, purchase orders, contracts, and GPO admin schedules. These are the most economically consequential records in the enterprise. They are also, by every independent measure, the least well governed.

Three observations drive the case:

1. **Vendor master data quality is materially worse than executives believe.** Independent assessments consistently show duplicate rates above 20%, parent–child relationships unmapped, and 30–60% of records carrying at least one defect.[1][2][3] Research published by MIT Sloan estimates poor data quality costs companies 15–25% of revenue.[4]
2. **Cybersecurity solved a structurally identical problem ten years ago.** Faced with controls that decayed silently between annual audits, the discipline moved to Continuous Controls Monitoring (CCM) — formally recognised by Gartner as an emerging GRC category, and now embedded in NIST, ISO 27001, COSO and COBIT practice.[5][6] Procurement and AP have not made the equivalent move.
3. **The cost of bad data has risen sharply.** Every finance and sourcing decision now being made on top of ERP data inherits whatever defects sit in the underlying records. A corrupted foundation produces wrong answers — and today those answers reach the board before anyone catches them.[7][8]

[1]GHX. "Item Master Data in Healthcare: Best Practices and Common Problems." 2024. https://www.ghx.com/the-healthcare-hub/item-master-data-healthcare-guide/ (accessed May 2026).
[2]TealBook (Alex Denomme). "Vendor Master Data Quality: Why Procurement Teams Struggle." 7 October 2025. https://www.tealbook.com/blog/vendor-master-data-quality-fix-the-foundation-before-you-build-anything-else/.
[3]GEP. "Vendor Master Data Hygiene: Why It Matters More Than Ever." 5 January 2026. https://www.gep.com/blog/technology/vendor-master-data-hygiene-attention-it-deserves.
[4]Fairmarkit. "Vendor Master Data," citing MIT Sloan research on data quality cost (15–25% of revenue; ~$3.1 trillion annually for the US economy). 2026. https://www.fairmarkit.com/glossary/vendor-master-data.
[5]Gartner. "Definition of Continuous Controls Monitoring (CCM)." Gartner Information Technology Glossary. https://www.gartner.com/en/information-technology/glossary/continuous-controls-monitoring-ccm (accessed May 2026).
[6]Quod Orbis. "Why Gartner thinks you need to know about continuous controls monitoring." 27 January 2026. https://www.quodorbis.com/a-simply-superior-approach-why-gartner-thinks-you-need-to-know-about-continuous-controls-monitoring/.
[7]ThoughtSpot. "AI Hallucinations and Biases: What They Are & How to Prevent Them." 31 July 2025. https://www.thoughtspot.com/data-trends/artificial-intelligence/ai-hallucinations-and-biases.
[8]Idera ER/Studio. "The Mirage of AI Hallucinations: How ER/Studio Fortifies Data Integrity." 11 March 2026. https://erstudio.com/blog/ai-hallucinations/.

THE ANALYST POSITION

This is not a back-office hygiene project. It is the data foundation on which every finance and supply-chain decision your organisation makes will stand or fall.

Providers that treat it as a one-time clean-up — or hand it to recovery firms who only get paid for what they find in the rearview mirror — will systematically trail the peers who make it a standing control.

The question stopped being whether. It is now who moves first — and whether they can do it in weeks rather than years.

Six practical conclusions follow.

1. **EBITDA terms, not savings rhetoric.** At single-digit margins, $1 recovered from supply-chain leakage is worth $20–$25 of new gross patient revenue. Reframe changes who owns the programme and prioritises project funding.
2. **Build the Golden Record before touching the analytics.** Every downstream tool — spend analysis, risk scoring, contract compliance — degrades to the quality of the underlying data. Build the foundation first.
3. **Use outcome-priced delivery for the first twelve months.** Derisk the program - The operator absorbs the implementation risk, not the buyer.
4. **Replace the contingency-recovery default with monthly continuous monitoring and management - the first cycle proves the economics.** The economics justify a phased approach.
5. **Name the control on the audit-committee dashboard.** Put continuous vendor-master quality alongside cyber and payor-mix. Until it is named, it will not be owned, measured, or funded.
6. **Run both tracks.** W-9 validation at onboarding stops new defects at the door. The Golden Record sweep cleans up what is already there. One without the other leaves a hole that compounds.

This paper sets out the evidence behind each claim, and what a defensible continuous-audit operating model looks like in practice.

1. The problem analysts keep finding, that executives keep underestimating

The market for procurement and AP audit services is splitting in two. Traditional firms — Big Four advisory, recovery-audit specialists, GPO consulting arms — sell engagements measured in months and millions. A newer generation of operators delivers the same underlying findings in days, without requiring ERP integration, priced on what they actually recover.[9]

The incumbents understand the threat. Accenture is in the middle of an $865M restructuring, exiting roughly 11,000 staff while retraining 550,000 in new tools; the Big Four are running equivalent programmes, with more than 100,000 automation-driven job cuts reported across consulting and adjacent industries in 2025 alone. But reskilling does not solve the structural problem. The cost base of a partner-and-junior pyramid is simply not the cost base of a software-driven operation running against the same data. That is the dilemma facing every incumbent: cannibalise the existing revenue model or defend it.[10]

1.1 The state of vendor master data in healthcare

The numbers on vendor master quality in healthcare are remarkably consistent — and remarkably bad. Gartner's procurement benchmarks, MIT Sloan's data-quality research, GHX's item-master studies, and the positioning material of the major recovery-audit firms all tell the same story, across three dimensions: how many suppliers a typical health system actually manages, how much money flows through that supplier estate, and what it costs when the underlying data is wrong.

[9]CNBC, citing Accenture CEO Julie Sweet. "Accenture plans on 'exiting' staff who can't be reskilled on AI amid restructuring strategy." 26 September 2025. Cited figures: ~11,000 positions cut under an $865M restructuring; 550,000 workers reskilled in generative AI fundamentals since 2023; AI/data professional headcount grown from 40,000 (2023) to 77,000 (2025). https://www.cnbc.com/2025/09/26/accenture-plans-on-exiting-staff-who-cant-be-reskilled-on-ai.html.

[10]Programs.com. "List of Companies Announcing AI-Driven Layoffs." Cited figures: over 100,000 employees impacted by AI-driven layoffs in 2025; 70,000+ in 2026 to date; 45+ CEOs have publicly attributed restructuring to AI efficiency gains. Updated 2026. https://programs.com/resources/ai-layoffs/.

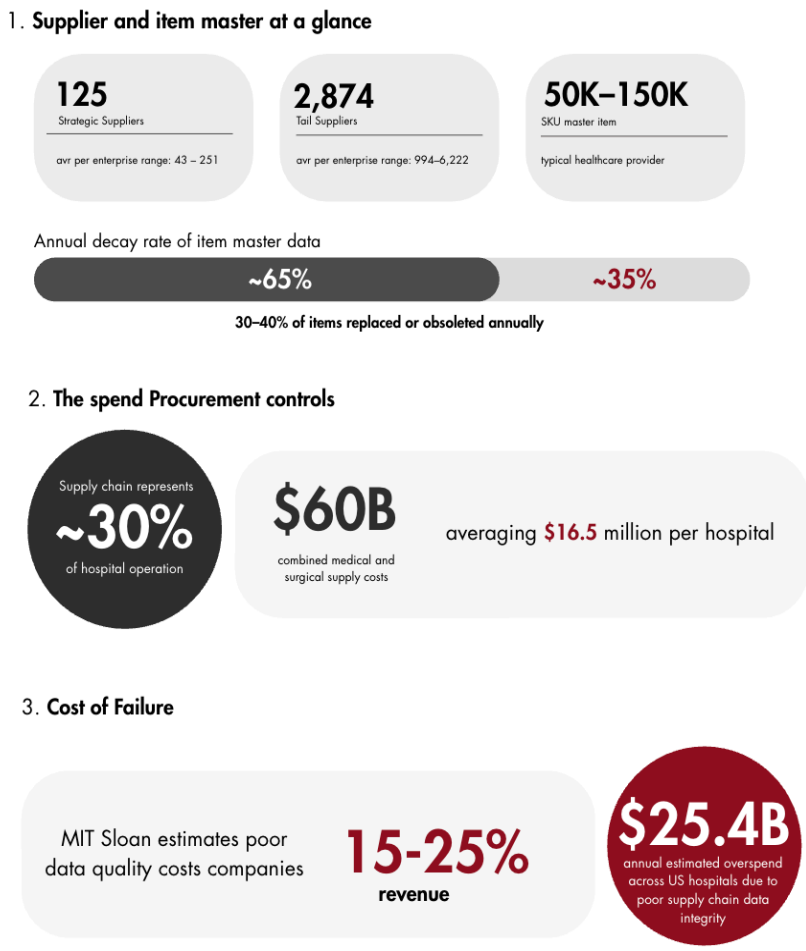

Three layers in one snapshot: the supplier and item-master population (Gartner Procurement Budget & Efficiency Benchmark; GHX item-master research); the spend procurement controls (~30% of hospital operating budget; ~$60B US medical and surgical supplies, averaging ~$16.5M per hospital); and the cost of failure (MIT Sloan 15–25% of revenue lost to poor data quality; ~$25.4B estimated annual overspend across US hospitals attributable to supply-chain data integrity gaps).

Each layer is well-evidenced, and each one tightens the argument:

- A typical enterprise carries a median of 125 strategically managed suppliers (range 43–251) and 2,874 tail-spend suppliers (range 994–6,222). For a mid-size US health system, layered on top of this is a hospital item master of 50,000–150,000 active SKUs, with 30–40% of items replaced or obsoleted in any given year.[11][12]

[11]Gartner. "Procurement Budget & Efficiency Benchmark." Gartner Information Technology research. Cited figures (median values, all industries): 125 strategic suppliers per enterprise (range 43–251); 2,874 tail suppliers per enterprise (range 994–6,222). https://www.gartner.com/en/documents/4526899.

[12]GHX. "Item Master Data in Healthcare: Best Practices and Common Problems" — cited figure: 30–40% of healthcare item-master data is replaced or obsoleted annually; typical hospital item masters range from 50,000–150,000 SKUs. 2024. https://www.ghx.com/the-healthcare-hub/item-master-data-healthcare-guide/.

- Research published by MIT Sloan estimates poor data quality costs companies 15–25% of revenue — approximately $3.1 trillion annually across the US economy.[13]
- In a representative VMCA (Vendor Master Continuous Audit) deduplication exercise on a US health system's AP and PO data, 11,853 raw vendor records resolved to 7,204 unique identities — a 23.6% duplicate-reduction rate — and only 36% of vendors appeared in both AP and PO systems. Roughly two-thirds of vendors lived in only one financial system of record.[14]
- Industry analysts estimate that bad data costs businesses trillions globally each year, with the average cost of a single data-driven error running into millions once churn and remediation are factored in.[15]

> *"Most procurement teams don't struggle because of process issues. They struggle because the vendor master doesn't reflect reality. The same supplier exists under multiple names. Remit-to addresses are outdated or missing. Classification codes are inconsistent. Parent-child relationships are unmapped."*
>
> — TealBook, 2025[16]

The defect is not occasional. It is not an edge case or an unlucky outlier. It is the baseline.

BUSINESS IMPACT — SPEND VISIBILITY

If you are intending to use AI within your business, which will impact or touch any financial parameter, if you don't solve your vendor master – hallucinations on your AI results will compound beyond your control.

Procurement agility research has found that high-performing organizations were 90% more likely to have high visibility into tier-1 suppliers; the gating factor was master-data quality, not analytics sophistication.[17]

[13]Fairmarkit. "Vendor Master Data," citing MIT Sloan research on data quality cost (15–25% of revenue; ~$3.1 trillion annually for the US economy). 2026. https://www.fairmarkit.com/glossary/vendor-master-data.

[14]hunterAI. Vendor Master Continuous Audit (VMCA) deduplication analysis on a US health system. AP data 2020-01 to 2026-02; PO data 2023-01 to 2025-05. Internal report, 2026.

[15]CX Today. "AI Hallucinations Start With Dirty Data: Governing Knowledge for RAG Agents." 23 November 2025. https://www.cxtoday.com/customer-analytics-intelligence/ai-hallucinations-start-with-dirty-data-governing-knowledge-for-rag-agents/.

[16]TealBook (Alex Denomme). "Vendor Master Data Quality: Why Procurement Teams Struggle." 7 October 2025. https://www.tealbook.com/blog/vendor-master-data-quality-fix-the-foundation-before-you-build-anything-else/.

[17]Deloitte, summarised in Fairmarkit. "Vendor Master Data" glossary, citing Deloitte research on procurement agility. 2026. https://www.fairmarkit.com/glossary/vendor-master-data.

1.2 Why the defect is structural, not operational

This is not a discipline problem solvable with better data-entry hygiene. Vendor master data degrades for predictable, structural reasons, and those reasons are amplified in healthcare specifically.

Domain	Symptom	Why it persists
Structural	Each region, business unit, or acquired entity creates suppliers independently. Formatting differs. Required fields vary. Validation steps get bypassed.	Healthcare providers are M&A-driven. Every acquired hospital arrives with its own vendor master, and the underlying systems are not merged immediately — acquirer and target typically operate parallel ERPs, AP and PO platforms for years until a full systems-integration program closes the gap. Until that happens, no single source of truth exists across AP, PO, contracts and GPO admin, and every consolidated spend number reported in between is reconciled by hand.
Technical	Procurement tools, finance platforms, ERP modules, and supplier portals store overlapping information. None sync perfectly. The GPO admin layer sees contract roster; the ERP sees AP records; neither sees the other.	AP-only and PO-only vendor populations diverge as soon as systems go live. Maintaining cross-system parity is no one's day job, and ERPs don't surface the gap until it appears as a control failure.
Business process	No one owns the clean-up cycle. Data decays silently between annual reviews. Vendor onboarding accountability sits across procurement, AP and category teams with no single point of failure.	Healthcare AP teams are chronically thin-staffed. Recovery audits run on 12–24 month contingency cycles and are scoped narrowly. Continuous quality assurance is structurally absent.
Cultural	Procurement is still measured on negotiated savings, not on the cleanliness of the data that proves savings. Data work is invisible work. Nobody gets promoted for fixing the vendor master.	Healthcare CPOs are increasingly accountable for purchased-services performance against the CFO — but the underlying data has never been brought up to a P&L-grade standard. The career incentive is to negotiate, not to clean.

Domain	Symptom	Why it persists
Design	Most ERPs treat vendor creation as a one-way write, not as a managed record with parent–child resolution, alias linking, or M&A awareness built in.	Healthcare ERPs (Workday, Oracle, Infor Lawson, SAP, Meditech) were architected before continuous data governance was an expectation. Bolt-on MDM products solve part of the problem but introduce another integration.
Economic	The cost of bad data is invisible at the line-item level and only materialises in aggregate. The CFO sees the leakage as part of the margin. No discrete budget owner is incentivised to fix it.	Recovery audit findings of 0.1–0.5% of AP are accepted as normal cost-of-business. Cyber treats a 0.1% control failure rate as unacceptable; procurement treats it as recoverable. The asymmetry is cultural, not analytical.
Technology	Traditional MDM solutions require pristine source data and 6–18 month implementations to deliver value. They are sold to data teams; procurement and AP cannot deploy them unilaterally.	The category is dominated by suites that demand integration before they produce insight. The cost of getting to the starting line is often higher than the cost of the unresolved problem looked at in isolation.

> *"Vendor data isn't static. It evolves over time, and if you don't monitor it, outdated records become the easiest entry point for fraud."*
>
> — Varun Kukreja, Zycus (in Trustpair, 2026)[18]

1.3 The downstream consequence — and why it is now urgent

For most of the last decade, the cost of unresolved vendor master data was absorbed into operating margin and recovered piecemeal by lookback audits. That equilibrium is ending. Three forces are breaking it simultaneously.

First, healthcare margins have compressed past the point where leakage can simply be absorbed. US providers run on single-digit or less operating margins. A 1% recovery on $1B of analysable spend delivers, in EBITDA terms, the equivalent of roughly $300M of

[18]Trustpair (citing Varun Kukreja, Zycus). "How to Manage Vendor Data: A Guide for Procurement." 6 February 2026. https://trustpair.com/blog/vendor-master-data-management-best-practices/.

new gross patient revenue. Supply-chain discipline is no longer a cost-reduction story. It is a revenue story — and it needs the same quality of data foundation.

Second, improper payments in the US healthcare system are quantified at scale. HHS estimated combined Medicare and Medicaid improper payments above $100 billion in fiscal 2023 — 43% of the government-wide total.[19] Provider-side AP and supply-chain leakage runs in parallel and is similarly under-controlled.

Third, every finance and procurement capability now being procured assumes a clean data foundation that does not actually exist. Layering software on top of a corrupted vendor master does not stabilise the data — it amplifies the defect, because the system averages across variants and returns authoritative-looking output that no one can easily disprove.

The result is wrong numbers dressed up as authoritative answers — reaching the board pack before anyone catches them. That is the shift most executives have not yet fully priced into their decision-making.

BUSINESS IMPACT — IMPROPER PAYMENTS AT SCALE

The Department of Health and Human Services estimated a combined total of over $100 billion in improper payments in Medicare and Medicaid programs in fiscal year 2023 — 43% of the government-wide total.

The mechanisms that generate errors (duplicate claims, coordination-of-benefits failures, eligibility errors, misapplied contract pricing) are the same mechanisms that generate AP leakage on the provider side.

The infrastructure to detect continuously, rather than recover them retrospectively, is the same.[20]

[19]US Government Accountability Office. "Medicare and Medicaid: Additional Actions Needed to Enhance Program Integrity and Save Billions." GAO-24-107487, April 2024. https://www.gao.gov/products/gao-24-107487.

[20]US Government Accountability Office. "Medicare and Medicaid: Additional Actions Needed to Enhance Program Integrity and Save Billions." GAO-24-107487, April 2024. https://www.gao.gov/products/gao-24-107487.

2. The data foundation problem — why dirty records are now a board-level risk

2.1 What modern analytics actually do with a corrupted vendor master

What happens when you feed a corrupted vendor master into a modern analytics system - The system does not flag the inconsistency. It averages across twelve variants of the same supplier and returns a fluent, plausibly-formatted answer. The user gets a number that looks right. It is not.

Industry commentary on this problem is now extensive:

> *"If your internal data is siloed, incomplete, or inconsistent, every AI initiative you launch will reflect and amplify those flaws. Think of it as: garbage in, hallucinations out."*
>
> — ThoughtSpot, 2025[21]

> *"When integrating your data with AI, AI will often 'polish' that bad data to make it look reliable and accurate."*
>
> — Idera ER/Studio, 2026[22]

Idera ER conclusion - when old-world dirty data at least looked wrong — a manager spotted the number, queried it, asked AP to investigate. Polished dirty data looks right. Boards approve it. Capital gets committed against it. The error only surfaces when a downstream control catches it, by which point the decision is already made and the money is spent or capital allocated.

BUSINESS IMPACT — DECISIONS MADE ON BAD DATA

Industry research has reported that approximately 47% of executives have admitted to making decisions based on faulty analytical output.

[21]ThoughtSpot. "AI Hallucinations and Biases: What They Are & How to Prevent Them." 31 July 2025. https://www.thoughtspot.com/data-trends/artificial-intelligence/ai-hallucinations-and-biases.
[22]Idera ER/Studio. "The Mirage of AI Hallucinations: How ER/Studio Fortifies Data Integrity." 11 March 2026. https://erstudio.com/blog/ai-hallucinations/.

In healthcare procurement specifically, an analytics tool asked to summarise spend by supplier from a vendor master with delta % duplication will produce category totals that are systematically understated, miscategorised, and unfit for negotiation.

The capital cost of acting on those numbers — in supplier-consolidation moves, in GPO renegotiation, in tier-1 risk reviews — runs into the tens of millions before the underlying defect is even identified.[23]

2.2 The structural fix is the same one cybersecurity adopted

This is not a new pattern. Cybersecurity faced an identical problem a decade ago: controls that looked correct when tested annually decayed silently between audits, and the cost of acting on stale assurance was catastrophic when it materialised. The discipline's response was to formalise Continuous Controls Monitoring (CCM) as a named category.

> *"Continuous controls monitoring (CCM) is a set of technologies to reduce business losses through continuous monitoring and reducing the cost of audits through continuous auditing of the controls in financial and other transactional applications."*
> — Gartner IT Glossary[24]

Two points in Gartner's framing matter for a healthcare procurement audience.

First, the definition explicitly names "financial and other transactional applications" as the canonical use case — not just cyber. CCM was always conceived as a discipline applicable to financial controls.

The discipline matured fastest in cybersecurity, as cyber had the largest acute incident-cost asymmetry – and the most visible impact in the media, if leaked, or hidden due to reputational brand management. Procurement has the same asymmetry now, on a larger absolute base.

Second, Gartner's own observation is that adoption lags the technology. "The concept and technology supporting Continuous Controls Monitoring is mature, but end-user

[23]Korra. "The $67 Billion Warning: How AI Hallucinations Hurt Enterprises (and How to Stop Them)." 22 August 2025. Cited figure: 47% of executives have admitted to making decisions on faulty AI content. https://korra.ai/the-67-billion-warning-how-ai-hallucinations-hurt-enterprises-and-how-to-stop-them/.
[24]Gartner. "Definition of Continuous Controls Monitoring (CCM)." Gartner Information Technology Glossary. https://www.gartner.com/en/information-technology/glossary/continuous-controls-monitoring-ccm (accessed May 2026).

adoption is not yet pervasive," the analyst firm noted in its hype-cycle commentary.[25] The barrier is operational maturity — the willingness to embed continuous control into the day-to-day operating model — not the existence of the tooling.

That barrier is exactly what an outcome-priced, integration-free service model is designed to lower.

2.3 The new nuance — purpose-built procurement intelligence

One structural shift was not yet in play when CCM first emerged in cybersecurity, and sharpens the requirement to address it – now.

The analytical tools now reaching healthcare finance are no longer general-purpose software trained on the whole internet. Gartner calls them Domain-Specific Language Models (DSLMs) — purpose-built systems trained specifically on the vocabulary, regulation and process logic of a defined field.

In procurement terms: tools that actually understand a GPO admin schedule, a 340B contract, or a clinical-preference SKU description, supported by factual evidence.

Gartner projects this market will reach $131 billion by 2035, with meaningfully higher reliability on domain-specific workflows than general-purpose tools.[26][27][28]

For healthcare procurement, this shift matters for three concrete reasons:

- A general-purpose tool trained on internet text cannot reliably interpret an item master, a GPO admin schedule, a vendor contract, a remit-to address variant, or a clinical-preference SKU or the opaque category of Purchase Services, etc. A procurement-specific system trained on these artefacts can. The accuracy

[25]Panaseer. "Gartner Hype Cycle: Continuous Controls Monitoring Category," summarising Gartner analyst commentary on CCM adoption maturity. 21 July 2020. https://panaseer.com/resources/blog/gartner-hype-cycle-continuous-controls-monitoring.

[26]Gartner (Roberta Cozza et al.). "Top Strategic Technology Trends for 2026: Domain-Specific Language Models." Gartner research, 2025. Cozza is a Senior Director Analyst on Gartner's Emerging Technologies and Trends team covering generative AI, multimodal AI, domain-specific models and agentic AI. https://www.gartner.com/en/documents/7013798.

[27]IBM, summarising Gartner research (Cozza, Searle; "AI Vendor Race: IBM Is the Company to Beat in Domain-Specific Language Model Enablement," 8 December 2025). "The next phase of the model AI race will not hinge on model raw capability or scale, but on how effectively the vendor can blend specific domain knowledge and reasoning, strong governance and deployment flexibility in a coherent ecosystem through DSLMs." https://www.ibm.com/new/announcements/ibm-recognized-by-gartner-as-the-company-to-beat-in-the-2025-ai-vendor-race.

[28]Gartner. "Domain-Specific Language Models: GenAI As a Precision Tool." Cited claim: DSLMs deliver up to 50% lower development costs versus generic LLMs, faster deployment and higher reliability for business-critical workflows; Gartner projects DSLM and DSLM-underpinned application market revenue will reach $131 billion in 2035. 2026. https://www.gartner.com/en/articles/domain-specific-language-models.

difference is not marginal; it is the difference between an output that survives audit and one that does not.

- These tools do not eliminate the data foundation problem — they intensify the return on solving it. The same system run against a clean Golden Record produces compounding value. Run against an unresolved vendor master, it produces faster, more confident errors.
- Automated decision chains built on these systems require clear accountability at every hand-off: validation, approval, supervision. Without a clean master record underneath, each step in that chain amplifies the errors from the step before rather than catching them.[29]

This is the key difference between the CCM moment in cybersecurity and the equivalent moment in procurement.

A decade ago, the tooling sitting on top of continuous controls was rule-based and predictable. Today it is statistical, domain-aware, and capable of acting without a human in every loop.

The lesson from cybersecurity holds: build the data foundation first, then put the tooling on top. But the cost of skipping that step is higher now than it was ten years ago, because the tooling is more powerful.

This shift — from generic tools to purpose-built procurement intelligence — is not an incremental upgrade. In consequence it is closer to the introduction of the smartphone than to a routine software release.

Healthcare procurement has to build its data foundation around what purpose-built procurement intelligence will demand of it — not around a recovery-audit script demanded a decade ago.

2.4 What cybersecurity did differently — and what healthcare finance and procurement should copy

Cybersecurity (2014–2020)	Healthcare (today)
Moved from annual SOC reports and quarterly self-attestation to continuous	Still relies on annual external audit plus 12–24 month contingency recovery sweeps.

[29]Gartner (Cozza et al., 2025). "Top Emerging Use Cases of Domain-Specific Language Models." Cited rationale: enterprises are augmenting or replacing general LLMs with DSLMs because generic models cannot meet performance, accuracy, compliance and relevance requirements for specialised enterprise workflows; multi-agent systems require accountable hand-offs of validation, approval and supervision across the chain of agents producing the output. https://www.gartner.com/en/documents/6569102.

Cybersecurity (2014–2020)	Healthcare (today)
control testing across NIST, ISO 27001, COBIT and COSO frameworks.	Vendor master is re-baselined on an irregular cycle, often only at ERP migration.
Built a named category — CCM — with Gartner, ISG and Forrester coverage. Created budget line items distinct from generic GRC.	No equivalent named category. Vendor master remediation sits inside ERP project budgets or one-off advisory engagements.
Adopted a controls-framework discipline: every control has a defined test, frequency, and owner. Decay is measured in days, not months.	Most controls (duplicate-payment prevention, off-contract leakage, GPO admin reconciliation) are tested transactionally, not as standing controls with measured decay rates.
Accepted that the human auditor is the bottleneck and built telemetry to feed analytics, with humans reserved for exception review.	Still scales with headcount. Recovery firms either charge contingency precisely because the work is labour-intensive, or quote a fixed price to rectify the problem as a one-time episodic event. Either model is a bandaid — the underlying defect is not fixed and resurfaces in the next cycle.
Treated data quality as a control objective in its own right. Bad data IS a control failure, not a precondition to one.	Treats data quality as a one-time data-cleansing project, separate from controls. Master-data hygiene is not on the audit committee dashboard.

THE ANALYST CONCLUSION

Healthcare finance & procurement is at the same moment cybersecurity was at in 2014 — but the precipitating event is sharper. The cost of acting on stale, ungoverned data was already acute. Domain-specific models intensify that cost further: they are a generational technology shift, not an incremental one.

The technology foundation is ready. The economic case is clear. The category will form — the only open question is whether the providers leading it will be the ones who defined it, or the ones who bought it later at a premium from whoever did.

3. The Golden Record — simple to state, hard to build, foundational to everything else

Every capability that matters in this domain — catching duplicate payments, closing off-contract leakage, reconciling GPO admin, scoring supplier risk, running clean spend analysis — depends on the same thing. You need to be able to look at any transaction and know, without ambiguity, exactly which supplier it belongs to: parent company, address, tax ID, NAICS classification, contract relationships and all. That resolved, deduplicated, continuously maintained identity – Industry term - the Golden Record.

Easy to state. Hard to build. And the gap between those two things is where the money is being lost right now — by providers and by the advisory firms they are paying to help them.

3.1 Why ERPs and MDM products do not solve this alone

Every healthcare provider's financial estate has three distinct layers, each carrying a different — and incomplete — picture of "the supplier." None of them can produce the Golden Record on its own. Example:

- The GPO admin layer sees provider finance data only as it is reconciled against one supplier across hundreds or thousands of provider members to generate a rebate payment schedule. Multiply that across thousands of suppliers, with individual contracts for products within a single supplier's portfolio, and the complexity is effectively unbounded. The admin layer cannot see off-contract vendors, alias variations, or payments made outside the GPO at all.
- The ERP / AP layer sees the vendor master, AP records, and PO records. It does not natively resolve parent–child relationships, M&A history, or alias variations across business units.
- The presentation layer — dashboards, finance reports, board packs — sees whatever the ERP exports, then runs it through layers of bespoke views that every organisation's management accounting function customises by hand for weekly, monthly, quarterly and end-of-year reporting.

  The result is a spider's web of manually-stitched data, sitting on top of the ERP, in which roughly 75% of finance team time is spent gathering and reconciling rather than analysing, and in which independent reviews have repeatedly found

that around 90% of spreadsheets contain errors and half of business-critical models carry material defects.[30][31]

The presentation layer cannot see what failed to reach the ERP from upstream, and the bespoke customisation amplifies rather than corrects the underlying defect.

No single layer can fix the gap from inside itself. Traditional services firms — Big Four advisory, recovery-audit specialists, single-vendor MDM products — all sell into one of these three layers. By definition, they cannot see across all three. The Golden Record needs something that operates between the layers, not inside any one of them.

3.2 What a resolved record actually contains

A defensible Golden Record validates a supplier identity across at least eight fields: parent company, category, address, city, state, zip, NAICS code, and TIN. The engine that produces it must do three things continuously, not once:

- Intelligent matching that recognises aliases, subsidiaries and naming variations across systems — not just exact-match deduplication.
- Context and validation against external sources (registry data, M&A activity, tax registrations) so the resolved record reflects current corporate reality.
- Closed-loop learning, so every audit cycle feeds the model and the next cycle finds defects faster.

BUSINESS IMPACT — WHAT A CLEAN GOLDEN RECORD UNLOCKS

In a recent VMCA exercise on a US health system (~$1.3B analysable spend), resolving 11,853 raw vendor records to 7,204 unique identities — a 23.6% duplicate reduction — was the precondition for four downstream capabilities: continuous duplicate-payment scanning, line-item anomaly detection at 98% precision on the resolved IDs, category-savings drift detection, and GPO admin-fee reconciliation against the actual paid spend. Without the resolution step, none of the four can be

[30]insightsoftware, citing CFO.com. "Can Excel Alone Ensure Accurate Financial Consolidation?" — cited figures: 63% of businesses globally rely on Excel for finance functions; 75% of finance teams' time is spent on non-value tasks like gathering data; only 22% of companies still rely exclusively on Excel for consolidation. 10 July 2025. https://insightsoftware.com/blog/can-excel-alone-ensure-accurate-financial-consolidation/.

[31]NextProcess, summarising Poon et al. (2024) review of spreadsheet error studies. "Why 94% Of Financial Spreadsheets Contain Errors (And What It Costs You)." Cited figures: ~90% of spreadsheets contain errors; half of spreadsheet models used in large businesses have material defects. 23 July 2025. https://www.nextprocess.com/capital-expense-management/why-94-of-financial-spreadsheets-contain-errors-and-what-it-costs-you/.

run reliably. With it, all four become recurring monthly outputs rather than annual recovery audits.[32]

3.3 Why this disrupts the traditional services model

Traditional procurement-advisory and recovery-audit firms have a specific shape: senior advisors scope, junior consultants extract and analyse, the engagement runs for months, the deliverable is a findings report, and the firm collects either billable hours or a cut of what they recover. That shape works when the underlying analytic work genuinely requires all those people and all that time.

Once the Golden Record exists and the right tooling sits on top of it, the labour intensity collapses.

A duplicate-payment sweep that once took six people six weeks now runs in 10 working days against your total supplier spend — no ERP integration, no on-site presence, a fraction of the contingency rate. The incumbents' cost structure simply does not survive that comparison.

> *"Modern AI platforms … don't require pristine master data as a starting point. Instead, they work with your existing messy data landscape and improve it incrementally."*
> — Suplari, 2025[33]

This is the architectural disruption. The traditional firm needs clean data to start. The outcome-oriented operator builds the clean record as the engagement, monetises the findings it produces, and leaves the provider with a recurring asset — not just a one-time report.

One caveat matters: the same tools that enable this disruption are also being purchased as point solutions inside the traditional services model. New tooling deployed against an unresolved vendor master — without structured validation built in, and without clear accountability across the output chain — produces faster errors, not fewer.

The architectural shift is not "add new tools to the existing engagement." It is "build a domain-aware control surface, with structured validation and human exception review

[32]hunterAI. Vendor Master Continuous Audit (VMCA) deduplication analysis on a US health system. AP data 2020-01 to 2026-02; PO data 2023-01 to 2025-05. Internal report, 2026.
[33]Suplari. "Procurement Data Quality — Why it's Important and How to Fix it." 2026. https://suplari.com/blog/procurement-data-quality.

built in, on top of a resolved record." That is the configuration that makes the labour-intensity collapse real rather than rhetorical.[34]

[34]Gartner (Cozza et al., 2025). "Top Emerging Use Cases of Domain-Specific Language Models." Cited rationale: enterprises are augmenting or replacing general LLMs with DSLMs because generic models cannot meet performance, accuracy, compliance and relevance requirements for specialised enterprise workflows; multi-agent systems require accountable hand-offs of validation, approval and supervision across the chain of agents producing the output. https://www.gartner.com/en/documents/6569102.

4. W-9 onboarding — the prospective control that catches what the lookback cannot

Continuous audit has two sides. The retrospective side resolves the vendor master that already exists — that is what most recovery-audit engagements try, imperfectly, to do. The prospective side stops new defects from entering the file in the first place. Most providers have no operating model for the prospective side at all. That gap is where the W-9 matters most.

4.1 Why the W-9 is the right control point

Form W-9, the IRS Request for Taxpayer Identification Number and Certification, is the regulatory document a US payer is required to obtain from any US vendor it will report on a Form 1099. The information it carries — legal name as filed with the IRS, federal tax classification, Taxpayer Identification Number (TIN), and a signed certification — is precisely the information needed to anchor a vendor record in a Golden Record schema. The form is not optional, it is not retroactive, and it is not subject to interpretation: it is a vendor-signed source document of legal identity.[35]

The IRS itself operates a free TIN Matching e-service for authorised payers, which validates name-and-TIN pairs against IRS records before payment is issued.[36] Best-practice guidance from across the AP-automation category converges on a consistent recommendation: collect W-9 forms during onboarding, before the first payment, and validate the TIN via the IRS TIN Match Program before the vendor record is activated.[37][38][39]

> *"Collecting W-9s during onboarding and verifying the information early eliminates most of the risk … Validating the information on the form prevents IRS mismatch notices, rejected filings, and the headaches that come with correcting errors after*

[35]Internal Revenue Service. "Instructions for the Requester of Form W-9 (03/2024)." https://www.irs.gov/instructions/iw9 (accessed May 2026).
[36]Internal Revenue Service. "Instructions for the Requester of Form W-9 (03/2024)." https://www.irs.gov/instructions/iw9 (accessed May 2026).
[37]Ramp. "W-9 Vendor Requirements Explained." 3 April 2026. https://ramp.com/blog/w9-vendors.
[38]Tipalti. "How to Use Form W-9 for Vendor and 1099 Compliance (2026)." 13 February 2026. https://tipalti.com/blog/w9-form/.
[39]Taina Technology. "Form W-9 Best Practices." https://www.taina.tech/resources-news-and-awards/form-w-9-best-practices (accessed May 2026).

> *you've already filed."*
> — Ramp, 2026[40]

The economic cost of getting this wrong is well documented. An incorrect TIN on a Form 1099 triggers an IRS CP2100 / CP2100A notice; if the issue is not resolved within 30 business days, the payer is legally obligated to apply 24% backup withholding to subsequent payments to that vendor, and is exposed to per-return penalties of up to $310 per incorrect form, with significantly higher penalties for intentional disregard.[41]

Tools that validate W-9 data in real time against IRS records have become standard in mid-market AP automation precisely because the cost asymmetry is so unfavourable: the validation is free, the consequence of skipping it is potentially unbounded.[42]

4.2 The volume problem in a mid-size hospital system

The often-overlooked feature of the W-9 control is volume. Gartner's Procurement Budget & Efficiency Benchmark, records a median of ~125 strategic suppliers per enterprise sitting on top of a tail population of ~2,874 suppliers (range 994–6,222), and a hospital item-master of 50,000–150,000 active SKUs of which 30–40% are replaced or obsoleted annually.[43][44] That is 600–1,200 new vendor records per year, every one of which becomes a permanent feature of the vendor master unless it is rigorously validated at the point of entry.

Three observations follow from this volume:

- **The W-9 is the cheapest defect-prevention dollar a provider can spend.** Validating a TIN against IRS records at onboarding costs essentially nothing. Resolving the same defect retrospectively — once it has produced misclassified spend, missed 1099 filings, or duplicate vendor entries — costs orders of magnitude more in time, audit fees, and IRS exposure.
- **Live-data W-9 collection is mostly redundant.** Existing vendors that are already being paid have, by definition, cleared whatever onboarding control was in place at the time. Re-collecting a W-9 from every active vendor is high-friction,

[40]Ramp. "W-9 Vendor Requirements Explained." 3 April 2026. https://ramp.com/blog/w9-vendors.
[41]TINCheck (Sovos). "Wrong TIN on a W-9 or 1099? Here's What to Do Next." 1 April 2025. https://tincheck.com/blog/wrong-tin-on-a-w-9-or-1099-heres-what-to-do-next/.
[42]TINCheck (Sovos). "Wrong TIN on a W-9 or 1099? Here's What to Do Next." 1 April 2025. https://tincheck.com/blog/wrong-tin-on-a-w-9-or-1099-heres-what-to-do-next/.
[43]Gartner. "Procurement Budget & Efficiency Benchmark." Gartner Information Technology research. Cited figures (median values, all industries): 125 strategic suppliers per enterprise (range 43–251); 2,874 tail suppliers per enterprise (range 994–6,222). https://www.gartner.com/en/documents/4526899.
[44]GHX. "Item Master Data in Healthcare: Best Practices and Common Problems" — cited figure: 30–40% of healthcare item-master data is replaced or obsoleted annually; typical hospital item masters range from 50,000–150,000 SKUs. 2024. https://www.ghx.com/the-healthcare-hub/item-master-data-healthcare-guide/.

low-yield work. The high-yield use of the W-9 control is on the inbound stream — new vendors, before first payment.

- **The retrospective view on existing vendors is a separate problem with a different solution.** For the existing population, the right control is identity resolution against external registers and TIN-matching of stored records — a sweep, not a form. The W-9 control is forward-looking. The retrospective control is the VMCA Golden Record exercise described in Section 3. Both are needed; neither replaces the other.

4.3 The two-track operating model

The two tracks are run as a single discipline by an embedded operator working an **As-an-Outcome (AaO)** engagement model — a short, scoped, outcome-priced engagement in which the analytic operator delivers a quantified, evidenced finding within a fixed window (typically 10 working days for the historical lookback), without ERP integration and without on-site headcount.

AaO replaces the multi-week-month, headcount-loaded recovery-audit engagement with a service-as-software primitive: the deliverable is the outcome, not the timesheet. Volume figures cited below (50–100 new vendors per month, 5,000–15,000 active vendor records) are consistent with Gartner's Procurement Budget & Efficiency Benchmark and GHX healthcare item-master research.[45][46]

Track 1 — Prospective (W-9 at onboarding)	Track 2 — Retrospective (Golden Record sweep)
Every new vendor: W-9 collected before first payment, TIN-matched against IRS records, name and address validated, parent–child relationship captured at onboarding.	Existing vendor population: dedupe across AP and PO systems, alias resolution, parent–child linking, external register validation, anomaly scan against payment history.
Volume: 50–100 new vendors per month for a mid-size health system. Manageable as a managed control inside the AP onboarding workflow.	Volume: 5,000–15,000 active vendor records in a typical mid-size health system. One-time sweep, then continuous monitoring as the maintenance layer.

[45]Gartner. "Procurement Budget & Efficiency Benchmark." Gartner Information Technology research. Cited figures (median values, all industries): 125 strategic suppliers per enterprise (range 43–251); 2,874 tail suppliers per enterprise (range 994–6,222). https://www.gartner.com/en/documents/4526899.

[46]GHX. "Item Master Data in Healthcare: Best Practices and Common Problems" — cited figure: 30–40% of healthcare item-master data is replaced or obsoleted annually; typical hospital item masters range from 50,000–150,000 SKUs. 2024. https://www.ghx.com/the-healthcare-hub/item-master-data-healthcare-guide/.

Track 1 — Prospective (W-9 at onboarding)	Track 2 — Retrospective (Golden Record sweep)
Cost asymmetry: validation is essentially free; non-validation creates IRS CP2100 exposure, backup-withholding obligations, and downstream master-data defects.	Cost asymmetry: every unresolved record degrades every downstream capability layered on top of it. The defect compounds.
Primary risk: IRS penalties, 24% backup withholding, 1099 mismatches, vendor relationship friction.	Primary risk: duplicate payments, off-contract leakage, GPO admin under-collection, wrong answers on aggregated spend data.
Owner: AP onboarding lead, supported by a structured W-9 collection and TIN-matching workflow.	Owner: CPO / CAE, supported by an AaO operator running the resolution sweep and the continuous monitoring layer.

BUSINESS IMPACT — WHY BOTH TRACKS ARE REQUIRED

Running only the retrospective sweep produces a clean Golden Record that immediately begins to decay as new vendors enter the file unchecked. Running only the W-9 prospective control produces a clean inbound stream that flows into a still-corrupted historical population. The two tracks are complementary, not substitutable.

The healthcare providers that operationalise both are the ones that can credibly report a continuously-managed vendor master to the audit committee — in the same way that cybersecurity teams now report continuously-tested controls rather than annual snapshots.[47]

[47]Taina Technology. "Form W-9 Best Practices." https://www.taina.tech/resources-news-and-awards/form-w-9-best-practices (accessed May 2026).

5. Vendor fatigue — why the traditional recovery-audit model is reaching the end of its useful life

Everyone talks about the contingency fee. Nobody talks about what recovery audits actually cost the AP team that has to process them.

The recovery-audit industry itself routinely acknowledges this dynamic in its own marketing: Premier Cost Recovery describes healthcare AP teams as "stretched thin" and positions its own approach as "low-disruption … without burdening your AP team,"[48] while Xelix — a category challenger — frames the same problem from the opposite side, observing that "AP teams must go back and forth with vendors to validate findings, gather evidence and negotiate repayments, all of which diverts resources from more strategic work."[49]

Run that across two or three audit firms simultaneously — which is common — and you get what practitioners call "vendor fatigue": a condition where the AP team spends a meaningful share of its working week chasing down findings that the firms themselves were not in a position to fully validate before handing them over.

5.1 The structural problem with the contingency-recovery model

Recovery audit firms are paid a percentage of what they recover. So their incentive is to surface findings — as many as possible — and hand them to you to validate. The result is predictable: the volume of findings handed over is far higher than the volume that will ever convert to cash, and your AP team absorbs every hour of validation work regardless of the outcome.

> *"AP teams must go back and forth with vendors to validate findings, gather evidence and negotiate repayments, all of which diverts resources from more strategic work."*
> — Xelix, 2026[50]

[48]Premier Cost Recovery. "Healthcare Accounts Payable Recovery Audits." Cited positioning: "Healthcare AP teams are stretched thin … our structured, low-disruption process identifies overpayments and delivers measurable recovery results without burdening your AP team." 2024. https://www.pcraudit.com/healthcare-accounts-payable-recovery-audits/.

[49]Xelix. "Why it's time to replace your AP recovery audit." 2026. https://xelix.com/recovery-audit (accessed May 2026).

[50]Xelix. "Why it's time to replace your AP recovery audit." 2026. https://xelix.com/recovery-audit (accessed May 2026).

Independent industry estimates put the typical recovery-audit firm yield at **0.1–0.15% of annual AP spend** — the equivalent of $1–1.5m on every $1bn of AP processed.[51] Set against an estimated total leakage rate of approximately $3.5m per $1bn of spend, this means the traditional model recovers, at best, a third of the leakage it is being paid to find.

The gap is not closed by running more audits; it is closed by moving the operating model from periodic, headcount-driven recovery to continuous prevention.

The validation problem is also acknowledged within the recovery industry itself. SpendMend, the largest dedicated healthcare AP recovery audit firm, publishes a client testimonial that is striking for its frankness about prior firms:

> *"Not only did failures in a previous audit firm's process result in lower recovery, but also some of the 'claimed' recovery dollars had to be paid back to vendors due to rejections."*
>
> — SpendMend client testimonial, 2024[52]

That is the false-positive problem at its most damaging. The findings were reported, the AP team chased them down, the recoveries were posted — and then reversed when vendors successfully contested them. Every cycle of that work is a tax on the AP function. Not in cash. In time and capacity.

5.2 The category itself is now acknowledging the limits of the model

In the last 18 months, multiple practitioner-side and category-side voices have converged on a consistent diagnosis, drawn from the public-facing commentary of Xelix, FiscalTec, Discover Dollar, and others.[53][54][55] The diagnosis is not that recovery audits are useless — they recover real money.

[51] Xelix. "Why it's time to replace your AP recovery audit" — cited figures: companies lose ~$3.5m per $1bn of spend to leakage; recovery audit firms typically recover only 0.1–0.15% of annual AP spend. 2026. https://xelix.com/recovery-audit.

[52] SpendMend. "AP Recovery Auditing." Published client testimonial referencing prior audit firm rejections. 3 February 2025. https://www.spendmend.com/ap-recovery-auditing/.

[53] Xelix. "Why it's time to replace your AP recovery audit." 2026. https://xelix.com/recovery-audit (accessed May 2026).

[54] FiscalTec. "Choosing the right Recovery Audit partner: A practical guide." 25 July 2025. https://fiscaltec.com/choosing-the-right-recovery-audit-partner-a-practical-guide/.

[55] Discover Dollar. "Why CPG Giants Use AP Recovery Audits to Cut Losses." 21 August 2025. https://www.discoverdollar.com/blog/why-cpg-giants-use-ap-recovery-audits-to-cut-losses.

The diagnosis is that they are the wrong primary operating model for the problem they are being asked to solve, because they are reactive, backward-looking, and structurally dependent on AP-team capacity to convert findings into cash.

> *"What started as a periodic diagnostic tool has, for many organisations, become a recurring dependency … errors slip through, a third party finds them months or years later, the business pays a commission to recover money that shouldn't have left in the first place and the underlying weaknesses remain. The same issues resurface in the next audit, and the loop continues."*
>
> — Xelix, 2026[56]

> *"A good provider lifts the burden, not adds to it with extra legwork and unnecessary demands."*
>
> — FiscalTec, 2025[57]

> *"Prioritize data hygiene to enable better matching and fewer false positives … Move from periodic to continuous audits to prevent future leakages."*
>
> — Discover Dollar, 2025[58]

Three observations converge in this commentary.

First, the false-positive rate of conventional recovery audits is high enough that practitioner-side commentators now explicitly recommend "fewer false positives" as a procurement criterion when selecting recovery partners.

Second, the cure for false positives is upstream data hygiene — i.e., a Golden Record.

Third, the operating model needs to move from periodic to continuous.

THE MAIN CONCLUSION

Each of these three observations independently points to the same architectural conclusion: the future of AP control is continuous, automated, prevention-oriented, and resting on a clean master record.

[56]Xelix. "Why it's time to replace your AP recovery audit." 2026. https://xelix.com/recovery-audit (accessed May 2026).
[57]FiscalTec. "Choosing the right Recovery Audit partner: A practical guide." 25 July 2025. https://fiscaltec.com/choosing-the-right-recovery-audit-partner-a-practical-guide/.
[58]Discover Dollar. "Why CPG Giants Use AP Recovery Audits to Cut Losses." 21 August 2025. https://www.discoverdollar.com/blog/why-cpg-giants-use-ap-recovery-audits-to-cut-losses.

The category's own incumbent voices are now stating, in their own marketing material, the case for the model that disrupts them.

5.3 The numbers: what changes when you move to continuous audit

The shift from recovery-firm-validated findings to system-validated findings is not incremental. It changes the time your AP team spends on validation by an order of magnitude.

Traditional recovery-audit cycle	Continuous-audit cycle
Engagement runs over 6–12 months. Recovery firm extracts data, runs proprietary scripts, surfaces findings, hands them to the AP team.	Historical sweep delivers evidenced findings in 10 working days. The system runs against the resolved Golden Record — dramatically lower false-positive rate from the start.
AP team validates each finding manually: pulls invoices, reconciles to PO and contract, contacts vendor, negotiates repayment, posts credit. Estimates from category practitioners place average statement reconciliation at ~30 minutes per supplier statement; complex statements take significantly longer.	The system pre-validates findings against the resolved record. Your AP team reviews exceptions only — not the full finding pipeline. Reconciliation that took 30 minutes per statement drops to 2–3 minutes.
AP team capacity displaced for the duration of the engagement. "Vendor fatigue" builds across multiple parallel audit firms each chasing the same supplier relationships.	Historical lookback: a matter of hours to weeks of AP-team capacity, not months. Ongoing continuous monitoring: approximately one hour per month of AP-team review on validated exception output.
Findings recovered retrospectively, often 12–18 months after the original payment. Some funds no longer recoverable. Some 'claimed' recoveries reverse on vendor challenge.	Findings surfaced in flight or within the current month. Vendor relationships stay intact because the resolution conversation happens close to the payment, not 18 months downstream.
Underlying defect is not fixed. The same issues resurface in the next audit cycle. Loop continues.	Findings feed back into the Golden Record and the controls baseline. Every cycle reduces the defect surface available to the next cycle.

The reconciliation time figures cited above are practitioner figures published by AP-automation specialists.[59] The order-of-magnitude reduction in AP-team validation time is the central economic argument for replacing periodic recovery audits with continuous monitoring — and it is the argument the recovery-audit industry itself is now being forced to address.

THE VENDOR-FATIGUE VERDICT

The traditional contingency-recovery model is not failing because the analytics are bad. It is failing because the operating model puts the burden of validation on a finance team that does not have the capacity to absorb it, recovers a fraction of the actual leakage, and leaves the underlying data defects in place to be re-discovered next cycle.

Continuous audit does not eliminate the recovery audit — it eliminates the conditions that make the recovery audit the only available option.

Hours for historical lookback. One hour a month for ongoing management. That is the operating model healthcare procurement will move to. The only question is which providers move first.

[59] FiscalTec. "The Cost of Errors in Statement Reconciliation: What AP Teams Must Know," citing IOFM research and the 2024 ACFE Report to the Nations. 17 December 2025. Cited figures: ~30 minutes per supplier statement manual reconciliation; 2–3 minutes with automation. https://fiscaltec.com/the-cost-of-errors-in-statement-reconciliation-what-ap-teams-must-know/.

6. Recommendations for healthcare boards, CFOs and CPOs

None of what follows requires a new technology budget. It requires a decision — to treat continuous audit of vendor and AP data as a named control with real governance behind it, the same way cyber controls are governed today. These are presented in the order a board or finance committee would most naturally encounter them.

Recommendation 1 — Anchor the program in EBITDA, not in savings rhetoric

Every dollar of recovered or avoided AP leakage is a dollar of EBITDA, booked at 1:1, with no payor-mix exposure, no A/R carry, no denial risk. At single-digit operating margins, that $1 is worth $20–$25 of new gross patient revenue. This is not a savings programme. It is a margin programme. Reframe who owns it, who funds it, and the seriousness.

Recommendation 2 — Start with the Golden Record, not with the analytics

Every analytic capability you will ever want to run — spend analysis, supplier risk, category benchmarking, sourcing automation, domain-specific procurement models — will produce output that is only as good as the record beneath it. The Golden Record is not a nice-to-have prerequisite. It is the leverage point. Build it first. Put the analytics on top. Reverse the order and your investment in analytics will deliver reliable results.

Recommendation 3 — Prefer outcome-priced delivery over headcount-priced delivery for the first 12 months

The risk is implementation overrun and value erosion in the gap between scoping and delivery — a gap that kills more internal programmes. For a capability yet to be proved, the choice is not about cost. It is about who absorbs the risk.

Recommendation 4 — Replace the contingency-recovery default with continuous audit

Traditional recovery-audit model places its highest cost on the AP team. Continuous audit reduces the historical lookback to hours and the ongoing monitoring to approximately one hour per month of AP-team review. The economics justify quick action.

Recommendation 5 — Name the control

Place "continuous vendor master and AP data quality" on the audit-committee dashboard, next to cyber controls, payor-mix monitoring, and clinical quality. As a named control with an owner, a metric, and a reporting cadence – will drive accountability and importance. The cybersecurity parallel is exact: CCM became a discipline the day boards started asking about it with the same regularity as breach exposure.

Recommendation 6 — Run both tracks: prospective W-9 and retrospective Golden Record

The W-9 control on new vendors prevents new defects from entering the file. The Golden Record sweep on existing vendors resolves the historical population. Continuous monitoring then maintains both. Skipping either track leaves a hole that compounds within months.

THE BOTTOM LINE

Healthcare procurement is at the same inflection point cybersecurity reached when CCM emerged. The data foundation problem is the same. The cost of ignoring it is now higher.

The technology works, the economics are compelling, and a new commercial model — outcome-priced, integration-free, measurable in days — removes the implementation barrier that has historically stalled these programmes.

If introducing AI into your organisation – this building block has to be corrected and managed continuously - failure to do so - all output associated with vendor spend will create "unseen" vulnerabilities and hallucinations that will jeopardise your business.

The case made in this paper is not a prediction. It is a description of what is already happening. The data exists. The operating model exists. The commercial structure that removes the implementation barrier exists.

What is not surfaced is the decision to treat vendor master and AP data quality as a named control, with the same governance weight as cyber and payor-mix. That decision is the one this paper is raising to healthcare boards, CFOs and CPOs to understand the risk at hand.

Methodology

All sources are footnoted where they appear. The goal is an honest read of where this market sits today — not a promotional document.

Source categories represented in the footnotes.

- Analyst and category research — Gartner (CCM, DSLM and Procurement Benchmark research; Roberta Cozza et al. on Domain-Specific Language Models), ISG, MIT Sloan, Panaseer, Quod Orbis, Deloitte (cited via secondary sources).
- Vendor master and procurement-data category — GHX, GEP, TealBook, Trustpair, Suplari, Fairmarkit.
- AP recovery audit category — SpendMend, Xelix, FiscalTec, Discover Dollar, Premier Cost Recovery.
- Data quality, model error, and DSLM commentary — ThoughtSpot, Idera ER/Studio, Korra, CX Today, IBM (citing Gartner).
- W-9, TIN matching and vendor onboarding compliance — IRS, Ramp, Tipalti, TINCheck (Sovos), Taina Technology.
- Improper payments and fraud in healthcare — US GAO, KFF, Withum, Healthcare Finance News.
- Services industry restructuring and automation-driven workforce shifts — CNBC (citing Accenture), Programs.com layoffs tracker, sector reporting on Big Four and management consulting.
- Finance reporting complexity and spreadsheet risk — insightsoftware (citing CFO.com), NextProcess (citing Poon et al. 2024 research review).
- hunterAI Vendor Master Continuous Audit (VMCA) deduplication analysis on a US health system, covering AP records January 2020 to February 2026 and PO records January 2023 to May 2025

Published by hunterAI as part of the Analyst Perspective Series. Written for CFOs, CPOs, Chief Audit Executives, and board members evaluating continuous-audit capability in healthcare provider settings. Share it with whoever would benefit from reading it. If something here is wrong, or if your experience differs, we want to hear about it.