AHI ARTIFICIAL HEALTHCARE INTELLIGENCE · 501(C)(3)

PART 2 OF 2 · WHITEPAPER

How One Document Can Fix American Healthcare

The Technology, the Evidence, and the Path Forward.

> *Part 1 traced the cascade from the encounter note to $5.3 trillion in broken healthcare. Part 2 examines why traditional reforms have struggled to fix it, why current technology in many cases makes the problem worse, and how AHI's MDMai and ENCOUNTERai are designed to reverse the cascade without requiring Congressional approval, regulatory reform, or EHR vendor cooperation.*

PUBLISHED BY
Artificial Healthcare Intelligence

DATE
February 2026

URL
artificialhealthcareintelligence.com

CONTENTS

Table of Contents

SUMMARY

Executive Summary

Part 1 of this white paper traced a ten-link causal cascade from the clinical encounter note to $5.3 trillion in healthcare spending with last-place outcomes among wealthy nations. That analysis identified the encounter note as the root cause, the single document through which all physician reimbursement flows, creating documentation burden that cascades through burnout, defensive medicine, declining reimbursement, access disparities, practice collapse, corporate consolidation, taxpayer-funded workarounds, emergency department misuse, industrialized claim denial, and socioeconomic patient sorting.

Part 2 answers the question that Part 1 inevitably provokes: Why don't we just fix it?

The short answer is that the U.S. healthcare system operates inside a political-economic structure that (through hundreds of millions of dollars in annual lobbying, regulatory capture, corporate profit extraction, political cycling, and deep institutional complexity) has proven extraordinarily resistant to the reforms that would address the cascade. Every dollar of the $5.3 trillion is someone's revenue. Every reform threatens someone's business model. Every legislative fix triggers a lobbying counter-offensive. Major healthcare reform has been defeated or substantially narrowed repeatedly over eight decades.

This white paper makes three main arguments:

First, traditional reform is structurally difficult. The healthcare lobby has spent over $10 billion in two decades influencing federal policy. Regulatory capture means the agencies tasked with implementing reform are often staffed from, and staffed into, the industries they regulate. Political cycling means that reforms that survive one administration are frequently dismantled or narrowed by the next. The ACA's Medicaid fee bump demonstrated that raising reimbursement rates improves access. Then 34 states let it expire. Similar patterns have recurred for decades.

Second, some current technology makes the problem worse, not better. The $35 billion HITECH investment built proprietary EHR silos that fragment patient data across many incompatible systems. AI ambient scribes, one of the fastest-adopted technologies in healthcare history, hallucinate clinical content at rates up to 31 percent per note in some studies, save only 34 to 41 seconds per note in randomized trials, and are, in the AAPC's words, "coding naive," with no demonstrated financial improvement in independent assessments. These tools automate a decades-old dictation workflow without addressing the billing architecture that created the documentation crisis.

Third, technology that operates directly at the encounter note, bypassing structural barriers, is a viable path forward. AHI's MDMai and ENCOUNTERai do not require Congressional approval. They cannot be filibustered. They do not depend on EHR vendor cooperation. They operate inside the existing billing architecture, within the existing regulatory framework, using existing EHR outputs. They target the document that generates downstream failures. If documentation burden drops from roughly 1,400 hours to closer to 100 hours per year, if undercoding is reduced, and if physicians reclaim their time and appropriate revenue, several parts of the cascade can begin reversing.

Fourth, the lobbying spend exists to protect a specific split: what the private system covers profitably versus what the public system absorbs at a loss. Part 1 documented how the encounter-note problem snowballs from the individual practice up through corporate consolidation at the state level and federal expenditure growth at the national level, while private payers preserve margin by offloading the chronically-ill and lower-income populations onto Medicaid, Medicare, FQHC grants, DSH, and EMTALA-funded emergency care. The federal government will spend roughly $2.4 trillion on healthcare in 2024 and roughly $4.3 trillion by 2033 on CBO projections. Chapter 8 of this paper lays out the seven specific policy asks the taxpayer should make in exchange for that $2.4 trillion: Medicare Advantage coding-intensity enforcement and overpayment recovery, Medicaid physician fee parity with Medicare, prior-authorization reform with teeth, denial and appeal transparency, narrow-network adverse selection oversight, public-program reimbursement parity with state-level cost increases, and a unified federal accounting of the workaround system's total cost. Deployment-first and policy-reform-next are not alternatives. They compound.

KEY POINT

The healthcare industry spent approximately $744 million in 2024 lobbying on federal policy. They cannot lobby to prevent a physician from using a free AI tool that saves time and recovers a meaningful share of their revenue. They cannot lobby to prevent the encounter note from being written correctly. The path around a lobbying wall is through the encounter note itself. That is what we built.

CHAPTER 1

WHY HEALTHCARE IS SO HARD TO FIX

The question every reader asks after seeing the evidence in Part 1 is: "Why don't we just fix it?" Why not raise Medicaid reimbursement? Why not fix EHRs? Why not stop the claim denials? The answer, after years of tracing every link in this chain, is that U.S. healthcare operates inside a political-economic system that, through lobbying, regulatory capture, corporate consolidation, political cycling, and institutional lock-in, resists exactly these fixes. The entities profiting from the $5.3 trillion system have spent decades constructing barriers to reform that have proven difficult for legislative action to overcome.

The $744 Million Lobbying Wall

The healthcare industry is the single largest lobbying sector in the United States, and it has held that position for over two decades. The health sector spent $743.9 million on federal lobbying in 2024, the only sector exceeding $700 million. The pharmaceuticals and health products industry alone spent $384.5 million, more than any single industry in any sector. The insurance industry added $155 million. Over the past 20 years, the health sector has spent approximately $10 to $12 billion on federal lobbying alone, not including state-level lobbying, campaign contributions, or dark money spending.

The spending is bipartisan by design. PhRMA, AHA, Blue Cross Blue Shield, and AMA all ranked in the top 10 overall lobbying spenders in 2024 and contribute to candidates on both sides of the aisle. UnitedHealth Group alone spent $9.93 million through the first three quarters of 2025, already exceeding its $7.52 million for all of 2024, according to OpenSecrets disclosures. News reporting has noted that large insurers' lobbying rosters routinely include former congressional staff, although the specific headcounts circulated publicly vary by source and year; readers should consult OpenSecrets registrant-level disclosures for current named staff.

Throughout 2024, the pharmaceutical, hospital, and insurance industries lobbied heavily against several reforms under debate. Some measures were narrowed, delayed, or dropped; others were ultimately included. Readers should consult the final December 2024 continuing resolution and subsequent reconciliation bills for the current status of PBM reform, site-neutral payment reform, and Medicare physician payment fixes, as those provisions have shifted repeatedly. The No Upcode Act, targeting Medicare Advantage overbilling, did not advance in the 2025 reconciliation process. Physicians received another year of Medicare payment cuts.

KEY POINT

Sources: OpenSecrets (2024–2025 lobbying totals); Axios, American Prospect, Star Tribune, Washington Post reporting on 2024 end-of-year legislative activity

The Private–Public Symbiosis: Why the Spending Is About Preserving the Split

Part 1 of this paper laid out the argument in detail. Private insurers, whose economics structurally disfavor enrolling lower-income and chronically ill patients, depend on government programs to absorb exactly those populations. Medicaid, CHIP, DSH payments, FQHC grants, Disproportionate Share Hospital funding, Medicare's traditional fee-for-service system, and the EMTALA-funded emergency department safety net together cover the people the private market is least profitable serving. This is not a moral judgment. It is a description of a stable equilibrium that every major actor has an economic incentive to preserve.

That equilibrium is what the lobbying spend defends. Every dollar spent to block Medicaid parity with Medicare defends the access gap that routes Medicaid patients into FQHCs at two to four times the private-practice per-visit cost, where the federal government pays. Every dollar spent to preserve the current Medicare Advantage risk-adjustment methodology defends roughly $23 billion in 2023 coding-intensity overpayments that federal taxpayers cover. Every dollar spent to block prior-authorization reform defends the denial machine that lets private insurers extract revenue retention at provider expense. The lobbying is not, principally, about blocking universal coverage. It is about preserving the specific split between what the private system covers profitably and what the public system absorbs at a loss.

Any reform strategy that does not understand this symbiosis underestimates the resistance it will face. A reform that would genuinely reduce private insurer revenue (universal Medicare expansion, a strong public option, national prior-authorization reform with teeth) will trigger coordinated industry opposition on a scale the last four decades of American politics have repeatedly demonstrated. A reform that would genuinely reduce federal healthcare spending by forcing the private system to cover the populations it now offloads would trigger the same response. This is why the most durable reforms in recent history have been the ones that increased total spending without altering the split. The equilibrium is resilient because every powerful actor profits from some part of it.

The implication for any serious reform proposal is unambiguous. If the goal is to bend the cascade this paper describes, the reform must operate at a point the lobbying apparatus cannot easily contest. That point is the encounter note. A provider cannot be lobbied out of using a tool that writes a better note. A biller cannot be lobbied out of filing an auto-generated appeal. An independent practice cannot be lobbied out of reopening its Medicaid panel when the arithmetic finally works. The political economy makes the encounter-note intervention attractive precisely because it is the one reform path that does not require the private-public split to be legislatively renegotiated.

KEY POINT

Sources: OpenSecrets federal lobbying data; MedPAC and HHS OIG on Medicare Advantage coding intensity; MACPAC supplemental payment and FQHC workaround analysis; Part 1 of this white paper on the wealth-health-insurance sorting machine

The Corporate Profit Motive: $5.3 Trillion Is Someone's Revenue

The fundamental economic reality is that every dollar of waste in U.S. healthcare is someone's income. The $812 billion in administrative costs (Himmelstein, Woolhandler & Campbell, using 2017 data), the $55.6 to $200-plus billion in medical liability system costs and defensive medicine, the $32 billion in avoidable ER visits, and the tens of billions of dollars in FQHC and supplemental payment workarounds all represent revenue streams for corporations with powerful incentives and resources to preserve them.

Industry-advocacy analyses indicate the seven largest health insurers' combined revenue roughly tripled in a decade, from approximately $511 billion in 2014 to $1.52 trillion in 2024, per EPIC for America's compilation of insurer annual reports. UnitedHealth Group's revenue reached approximately $400 billion in 2024, roughly 4.6 times its 2009 level of about $87 billion. During the same decade, physician Medicare payment fell approximately 33 percent in real terms. Peer-reviewed and policy research has documented trillions of dollars in cumulative shareholder payouts (dividends plus stock buybacks) from publicly traded U.S. healthcare companies over the past two decades, with annual totals in the range of $150 to $200 billion in recent years, during a period of rising physician burnout, declining access, and worsening outcomes.

The Partnership for America's Health Care Future (PAHCF), an alliance of hospital, insurance, and pharmaceutical interests, formed in 2018 with a stated objective of preventing single-payer healthcare, Medicare expansion, or a public option. Leaked planning documents describe objectives to "change the national conversation around single payer / Medicare for All" and "minimize the potential for this option from becoming part of a national political party's platform in 2020." Members spent a combined $143 million on lobbying in 2018 alone.

KEY POINT

Sources: EPIC for America (2024, secondary compilation of insurer annual reports); peer-reviewed and policy research on healthcare shareholder payouts; The Intercept; Jacobin (2024)

Regulatory Capture: The Regulators Often Come From and Return to the Regulated

A Health Affairs study (2023) documented the revolving door at HHS. Among people appointed between 2004 and 2020, roughly 15 percent came directly from private industry and 32 percent exited to industry after their tenure. The greatest net exits were from the CDC and CMS, the two agencies most central to healthcare

regulation. More than half of appointees at CMS, CDC, and the Office of the Deputy Secretary went to private industry after government employment. CMS Administrator Seema Verma went to Oracle. CMS Administrator Marilyn Tavenner came from 25 years at for-profit hospital giant HCA. The pattern is bipartisan and persistent.

The AMA's Relative Value Scale Update Committee (RUC), which effectively sets Medicare physician payment rates, has been described as perhaps the most consequential example of regulatory capture in healthcare. CMS accepts RUC recommendations roughly 90 percent of the time, and the committee has systematically favored procedure-heavy specialties over primary care, contributing to the primary care crisis that drives the cascade.

KEY POINT

Sources: Health Affairs (2023); Healthcare Dive (2023); Milbank Memorial Fund (2022)

Political Cycling: Healthcare Reform's Short Half-Life

The history of healthcare reform includes many ambitious proposals that were killed or narrowed by similar forces: Truman's universal health insurance (1948, opposed by the AMA), Clinton's Health Security Act (1993, defeated after a sustained industry-wide opposition campaign widely estimated to have cost hundreds of millions of dollars across industry coalitions, including the Health Insurance Association of America's "Harry and Louise" advertising), the ACA's public option (2009, dropped amid intense insurance industry pressure), and Medicare for All (2019–2020, opposed by PAHCF despite majority public support in polls and more than 100 Congressional co-sponsors).

The most instructive case is the ACA's Medicaid primary care fee bump. In 2013 and 2014, the ACA temporarily raised Medicaid primary care rates to Medicare parity. The evidence was unambiguous: appointment availability increased by 7.7 percentage points, and each $10 increase in fees produced a 1.7 percentage point increase in appointment probability. Then 34 states let the increase expire in January 2015, triggering an average 42.8 percent reduction in primary care fees. The policy demonstrated that the fix worked. Then it was allowed to expire.

KEY POINT

Sources: PMC / Obama JAMA article (2016); KFF (2025); Urban Institute; Polsky et al., NEJM (2015)

The Complexity Barrier

Perhaps the most insidious barrier is epistemological. The healthcare system is so complex that most people, including many policymakers, have difficulty identifying the root cause. Current political debates focus on drug pricing (approximately 10 percent of total spending), PBM reform (CBO estimates savings of less than 0.1 percent of premiums), and hospital consolidation (real but downstream of practice economics). None addresses the encounter note.

The CPT coding system contains approximately 11,000 codes with 420 annual updates. The billing and documentation requirements are so complex that administrative errors (including incorrect coding, missing information, and eligibility issues) account for over 80 percent of preventable claim denials across specialties. This complexity is not an unfortunate byproduct. It functions as a mechanism through which administrative costs flow disproportionately to providers while payers automate denials at roughly 1.2 seconds per claim.

KEY POINT

When someone says “just raise Medicaid reimbursement,” they are proposing a policy that must survive the Senate filibuster (60 votes), overcome hospital, FQHC, and insurer lobbying, sustain funding across multiple administrations, and survive CMS implementation by revolving-door officials. When someone says “just fix the EHRs,” they are proposing replacing $35 billion in HITECH-subsidized infrastructure against vendor lobbying. When someone says “just stop the claim denials,” they are proposing reducing insurer revenue by billions against an industry that has successfully blocked prior authorization reform for decades. Each “simple fix” triggers opposition from multiple powerful industries at once.

KEY POINT

Sources: Health Affairs; CAQH Administrative Cost Report; CBO scoring data

CHAPTER 2

WHY EHRs ARE HARD TO FIX AND AI ALONE WILL NOT SAVE THEM

The $35 Billion Subsidy That Built Proprietary Silos

The HITECH Act of 2009 allocated approximately $35 to $38 billion in incentive payments to drive EHR adoption. Physician adoption jumped from 42 percent to 88 percent, and hospital adoption climbed from 9 percent to 96 percent. But the program's fatal design flaw was mandating that providers digitize records without specifying how systems should be built. No unified architecture. No enforceable data standards. No interoperability requirements. Billions in public funds effectively subsidized the entrenchment of proprietary, incompatible systems.

An HHS Office of Inspector General audit (2017) raised concerns about substantial improper Meaningful Use incentive payments to professionals who did not clearly meet program requirements, with estimated exposure in the hundreds of millions of dollars; readers should consult the published OIG audit for exact figures. Academic research has also suggested that the HITECH subsidies accelerated EHR adoption by only a modest margin above what market forces would have achieved on their own.

- **The Repository That Was Never Mandated**

The deeper design failure was structural. HITECH did not mandate a single longitudinal repository of patient records. It did not require EHR vendors to interoperate. It did not specify who would hold a patient's unified chart, how to address the same patient across systems, or what must happen when a patient crossed from one organization's EHR to another's. The consensus decision at the time was to let competitive markets produce interoperability. They did not. It took the 21st Century Cures Act of 2016, and its final rule in 2020, for the federal government to formally name information blocking as something providers and vendors could not do, and enforcement has lagged well behind complaints filed.

The operating condition this produced is not abstract. Consider a clinical arrangement that is typical rather than exotic. A private clinic runs its own EHR. The same providers round at two local hospitals that are both owned by the same Healthcare company and both run Cerner, but each hospital maintains its own siloed Cerner instance that does not share data with the other. The same providers also see patients at a partnered FQHC that runs a different EHR again. Four systems. The same patient can appear in all four. A treating clinician at any one site, at the point of care, sees only the fraction of that patient's record that lives in the system they happen to be logged into. There is no unified place to go look.

Walk a single patient through the workflow and the danger becomes concrete. The patient is seen in the clinic and told to go to the ED. They are admitted by the hospitalist at Hospital A, whose EHR instance does not automatically pull in the clinic note. The patient clinically deteriorates and is transferred to Hospital B for a higher level of care. The transferring hospital is responsible for sending the chart with the patient. Sometimes it arrives. Often it does not, or arrives only partially, or arrives as a fax the receiving clinician will not read because the patient is already in front of them with an urgent presentation. The receiving clinician, in practice, is working from what the patient can tell them, which has been measured in peer-reviewed studies as incomplete or inaccurate in a material share of cases. This is not a corner-case of American medicine. It is a standard operating condition.

Interoperability Remains a Promise, Not a Reality

Despite more than a decade and a half of regulatory effort, ONC survey data has repeatedly found that a minority of independent hospitals exchange data routinely, and that clinicians rate outside-organization integration as one of the weakest aspects of their EHR experience. Patient data fragmentation is severe: Medicare beneficiaries' records are commonly spread across multiple EHR vendors over a patient's lifetime, and patient matching accuracy drops sharply across organizations (widely reported at roughly 50 percent in cross-system exchanges). TEFCA took seven years from authorization to launch and remains voluntary. Enforcement actions for information blocking have lagged far behind complaints filed.

ONC Data Brief 71 (May 2024) reported that only 42 percent of non-federal acute care hospitals routinely send patient health information to other hospitals, and only 38 percent routinely send to ambulatory providers outside their own organization. Those are the rates more than a decade into HITECH and four years into the Cures Act final rule. The other side of the exchange matters just as much: even when a record is sent, the receiving clinician has to find it, trust the patient match, and read it in time to act on it.

- **What Fragmentation Costs in Lives and Limbs**

The safety literature on fragmented records is consistent and difficult to read dispassionately. The Joint Commission has repeatedly attributed roughly 80 percent of serious medical errors, at least in part, to communication failures during handoffs. AHRQ's Patient Safety Network has cited figures in the same range, with analyses pointing to communication breakdowns during transitions of care as a contributing factor in a large majority of preventable adverse events. Makary and Daniel estimated in BMJ (2016) that medical error was associated with roughly 251,000 deaths per year in the United States; later analyses have debated the point estimate but not the underlying finding that the count is large enough to place medical error among the leading causes of death.

Interhospital transfer is the sharp end of the problem. A frequently cited Critical Care Medicine study examined patients transferred into tertiary ICUs and reported an adverse event rate of approximately 42 percent within the first 24 hours of arrival, in-hospital mortality in the range of 17 percent for the transferred cohort, documentation completeness below 60 percent for records arriving with the patient, and missing clinical information in roughly one-third of transfers. Medication reconciliation studies have documented error rates above 60 percent on admission, with undocumented allergies and omitted home medications being the most common categories. Patient self-report, which is what the receiving clinician falls back on when the chart does not arrive, has been measured in peer-

reviewed work at roughly 81 percent agreement with the record for medications and with allergy-list discrepancies approaching 41 percent. This is not dishonesty. It is memory, stress, language, cognitive status, and the simple difficulty of listing every medication a complex patient takes.

Malpractice data is the downstream mirror. CRICO Strategies reported in 2015 that approximately 30 percent of malpractice claims involved communication failures, with total associated economic impact estimated at $1.7 billion and an associated death toll counted in the thousands. Candello's more recent analyses, through 2024–2025, have reported the communication-failure share of claims closer to 40 percent, suggesting the problem has not improved and may be worsening as care becomes more distributed and time-pressured.

None of these statistics live only on a dashboard. They are what the workflow described above produces, every day, in every market where patients legitimately receive care across more than one organization.

Epic's Inescapable Gravity

Epic controls a plurality of U.S. acute care hospitals (approximately 42 percent by bed market share, per KLAS) and generated an estimated $5.7 billion in revenue in 2024. Large implementations routinely exceed $1 billion, including Mass General Brigham at approximately $1.2 billion and Mayo Clinic at approximately $1.5 billion (figures per public reporting from the respective health systems). Antitrust complaints filed in 2025 against Epic describe the EHR platform as one that is, in practice, extremely difficult for hospital customers to leave once implemented, reflecting high switching costs rather than a genuinely competitive marketplace. Physician satisfaction scores on EHR usability remain low in many surveys, yet customer retention remains very high because migration away from an entrenched EHR is prohibitively expensive.

The VA's Oracle / Cerner implementation is a definitive cautionary tale. VA Office of Inspector General reports, Government Accountability Office analyses, and media coverage have estimated total program costs on the order of $37 billion or more over roughly 28 years, with some external projections reaching $50 billion when lifecycle costs are included; exact figures vary by source and assumption. Only a small fraction of VA medical centers had gone live after several years of deployment and approximately $5 billion in spending.

The Note Is the Real Problem, Not the EHR

U.S. clinical notes are several times longer than notes in other English-speaking countries using the same EHR software. U.S. clinicians spend 90.2 minutes per day actively using the EHR, versus 59.1 minutes for non-U.S. clinicians (approximately 53 percent more time) using the same systems. This suggests the dysfunction is not a pure technology problem. It is a documentation-for-reimbursement problem rooted in the U.S. payment architecture.

Wang et al. found that only 18 percent of progress note text was manually entered by physicians, with 46 percent copied from prior notes and 36 percent imported from other structured sources. Rule et al. found median outpatient note redundancy reached 58.8 percent by 2018. AI ambient scribes, in their current form, largely automate the data entry layer without challenging the billing architecture that demands voluminous documentation.

KEY POINT

Sources: HITECH Act data; 21st Century Cures Act (2016) and Final Rule (2020); ONC (2023); ONC Data Brief 71 (May 2024); KLAS Research; Epic Research (2023); VA OIG and GAO reports; Joint Commission on handoff communication; AHRQ Patient Safety Network; Makary and Daniel, BMJ (2016); Critical Care Medicine interhospital-transfer study (2016); medication reconciliation literature; CRICO Strategies (2015) and Candello (2024–2025); Sinsky et al.; Rule et al., JAMA Network Open (2021)

CHAPTER 3

THE CASE AGAINST AI AMBIENT SCRIBES

AI ambient scribes, among the fastest-adopted technologies in healthcare history, rest on a premise that we believe is flawed: that converting spoken words to written notes solves the documentation crisis. The evidence from a growing body of peer-reviewed studies suggests otherwise.

Hallucinations, Omissions, and Fabrications

In one evaluation, 31 percent of AI-generated ambient notes contained at least one detected hallucination versus 20 percent of physician-written notes (Palm et al., Frontiers in Artificial Intelligence, 2025). A large manual evaluation (Asgari et al., npj Digital Medicine, 2025) found a hallucination rate of 1.47 percent at the sentence level, with 44 percent classified as "major," meaning they could impact patient diagnosis or management. The most dangerous hallucinations appeared in the Plan section (21 percent), the section containing direct instructions for patient care.

An OHSU study using ChatGPT-4 found an average of 23.6 errors per clinical case. Only 52.9 percent of data elements were correctly reported across replicates of the same encounter, suggesting that AI scribe output can be fundamentally unreproducible. OpenAI's Whisper model, used by approximately 30,000 clinicians, was found in one analysis to fabricate sentences in approximately 1.4 percent of transcriptions, including violent and racially charged fabrications triggered by speech pauses.

The Editing Burden: Trading Typing for Proofreading

A UCI Health study of 23,760 AI-drafted notes found that 84.4 percent were edited by clinicians before signing. The first randomized clinical trial of AI scribes (Lukac et al., NEJM AI, 2025) found users saved an average of just 41 seconds per note (approximately a 9.5 percent reduction), while DAX Copilot savings did not reach statistical significance. Kaiser Permanente's 63–week deployment, the largest real-world evaluation at 7,260 physicians and 2.5 million encounters, reported that among non-users, barriers included "the perception that editing AI-generated notes took more time than typing from scratch."

Physician satisfaction with AI scribes has been mixed in reported deployments, with health systems reporting roughly balanced shares of enthusiastic adopters and frustrated non-users. Readers should look to each institution's own published data for specific Net Promoter Scores and retention rates.

A 30-Year-Old Workflow in New Packaging

The concept of converting spoken words to written clinical documentation goes back more than a century. Whether the intermediary is a stenographer (1900s), a transcriptionist (1960s-1990s), Dragon NaturallySpeaking (1997), or a GPT-4 ambient listener (2023), the workflow is largely unchanged: the provider speaks, a machine converts speech to text, the provider reviews the output. Microsoft acquired Nuance for $19.7 billion in 2022, and DAX Copilot is marketed as "the latest addition to the Dragon family."

The core limitation is irreducible. Ambient scribes can only document what is spoken aloud. Clinicians who have piloted these tools have repeatedly noted that physical-exam findings and cognitive clinical reasoning often go uncaptured unless the physician narrates them out loud. The Assessment and Plan section, the most clinically important part of the note, is systematically under-documented by pure transcription because clinical reasoning is cognitive work that rarely gets spoken during the visit.

Coding Naive and Billing Blind

The AAPC concluded that most AI scribes are "coding naive, unaware of all the requirements of documentation from a revenue cycle perspective." The Peterson Health Technology Institute's March 2025 independent assessment found that health systems "haven't seen significant financial improvements" from scribe technologies. Reports in the field have also raised concerns that generative AI scribing, absent billing optimization, may increase note length rather than reduce bloat. Institutions evaluating AI scribes should assess note length changes directly in their own deployments.

KEY POINT

AI ambient scribes today represent an incremental improvement on a decades-old dictation workflow. They still require the provider to verbalize everything, still generate errors requiring physician review, still often produce verbose notes, still struggle in noisy or complex clinical environments, still largely ignore billing optimization, and still lack FDA oversight while placing full legal liability on the signing physician.

KEY POINT

Sources: Palm et al. (2025); Asgari et al. (2025); Lukac et al., NEJM AI (2025); Kaiser Permanente, NEJM Catalyst (2025); AAPC (2024); Peterson Health Technology Institute (2025)

CHAPTER 4

THE TECHNOLOGY THAT FIXES THE NOTE

If the encounter note is the root cause of healthcare's dysfunction, and if traditional reform is structurally difficult, and if current technology largely automates the problem rather than solving it, then the solution must be a technology that operates directly on the encounter note, within the existing system, without requiring permission from any of the entities profiting from the status quo.

AHI has built two interlocking platforms designed for exactly this purpose. MDMai (Medical Decision Making AI) analyzes existing encounter notes for billing accuracy and documentation quality. ENCOUNTERai generates complete encounter notes predictively from minimal structured input, trained on each provider's own clinical documentation patterns.

MDMai: What It Does and Why Nothing Else Does It This Way

MDMai is designed around a single objective: to make sure every encounter note it processes is billed at the most accurate and most defensible level the documentation will support. It is not a scribe, a claims scrubber, or a clearinghouse. It sits between the signed encounter note and the superbill, in the narrow window where the clinical record has already been created but billing decisions have not yet been made. That position is the one place in the revenue cycle where documentation can still drive coding rather than the other way around.

The problem MDMai is designed to solve has two compounding structural drivers. The first is that many notes are not signed at peak clinician attention. Arndt et al. (2017, Annals of Family Medicine) found primary care physicians spend an additional 86 minutes per day on EHR work after clinic hours, the figure widely cited as "pajama time," and subsequent JGIM and JAMA analyses have confirmed the pattern across specialties. A note signed at ten o'clock at night by a physician who has already seen 25 patients is rarely a note whose every billable element has been fully refined. The clinician wants to finish, sign, and go to bed. The downstream consequences of that understandable decision are significant, and they almost always favor the payer.

The second structural driver is that the U.S. medical coding system has grown into a thicket no single clinician or biller can be expected to hold in working memory. CPT alone spans thousands of codes. ICD-10-CM contains roughly 70,000. HCPCS Level II adds several thousand more. Each set ships with sub-rules, modifiers, specialty carveouts, payer-specific overlays, and per-code documentation requirements. The reference libraries used by full-time trained coders literally fill shelves. Two concrete examples show how that complexity translates into invisible revenue loss at the note level.

The first example is sepsis specificity. ICD-10-CM distinguishes between plain sepsis (A41.x), severe sepsis without septic shock (R65.20), and severe sepsis with septic shock (R65.21). Plain sepsis is a Complication or Comorbidity (CC) under MS-DRG grouping. Severe sepsis and septic shock both carry Major Complication or Comorbidity (MCC) status, which can move the case into a meaningfully higher-weighted DRG family (DRG 871 versus DRG 872, for example). A physician who writes "sepsis" in a note where the clinical data actually support acute organ dysfunction, and therefore severe sepsis, has unintentionally downcoded their own documentation. The note will be coded at the less specific level because that is what it literally says, and legitimate DRG weight can walk out the door on a single admission. AHIMA, AAPC, and ACDIS have documented this pattern in clinical documentation integrity practice for years. The same under-specificity dynamic plays out in pneumonia (organism-specific versus unspecified), diabetes (with versus without named complications), heart failure (staged versus unstaged), and many others.

The second example is critical care. CPT 99291 and 99292 are criteria-based codes, not location-based codes. The governing criteria are a patient with critical illness or injury having a high probability of imminent or life-threatening deterioration, direct high-complexity decision-making by a physician or qualified healthcare professional, and total critical care time on the date of service of at least 30 minutes for 99291, with each additional 30 minutes captured by 99292. Nothing in the CPT or CMS definition restricts the code to an ICU or a critical care ward. Critical care can be billed in the emergency department, a step-down unit, a telemetry floor, a regular medical-surgical floor, or during a rapid response or code event, as long as the clinical criteria and time thresholds are met. (CMS E/M Services Guide, September 2024; AMA CPT Professional Edition.) Most clinicians who have not been trained as billers assume the opposite. Because they assume the opposite, they do not document the critical care time they actually provided, and the code never gets filed.

These are two examples of many. The point is not that providers are bad clinicians. The point is that the coding system rewards full-time trained billers and penalizes everyone else, and most independent practices cannot afford full-time trained billers. One could almost read the shelf-filling reference libraries and their overlapping rule sets as a system engineered to ensure no single practitioner could master the whole thing. MDMai is built to close that gap at the note level, before the superbill is created.

- **The Workflow MDMai Applies to Every Note**

MDMai is a multi-agent AI system. When a signed encounter note enters the pipeline, it proceeds through five stages.

Stage One: Path Selection. For office and outpatient E/M visits (99202 to 99215), the 2021 E/M revisions adopted by AMA and CMS (effective January 1, 2021 and still in effect) allow the provider to select the billing level based on either total time on the date of encounter or medical decision-making complexity. Both paths remain valid, and they often yield different codes. MDMai evaluates each path against the documentation and recommends whichever yields the higher defensible code. This is not upcoding. It is the allowable election the CPT rules created specifically so that clinicians would not be penalized for spending more time with complex patients than the MDM structure alone rewards.

Stage Two: Time Reconstruction and Prolonged Services. When the time-based path is preferred, MDMai applies algorithms built specifically to reconstruct total clinician time on the date of service (face-to-face plus non-face-to-face time, as the 2021 rules define it) from the clinical evidence in the note. If the reconstructed time crosses the prolonged-services threshold, the appropriate add-on code is identified: CPT 99417 for commercial payers, HCPCS G2212 for Medicare, and the analogous inpatient codes (CPT 99418 commercial, HCPCS G0316 Medicare inpatient, G0317 skilled nursing). Prolonged services are among the most commonly missed add-ons in outpatient E/M, and the revenue impact is material.

Stage Three: Billable-Element Refinement. Before any billing codes are generated, MDMai refines every billable element against the clinical data in the note. Diagnoses are pushed to the most specific ICD-10-CM level the documentation actually supports (the sepsis case is one instance; under-specificity in pneumonia, diabetes with complications, and heart failure staging are others). Procedures are checked for the most specific CPT code and correct modifiers. HCPCS Level II codes are scanned for separately billable supplies, drugs, and services. Critical care criteria are evaluated independently of encounter location. Chronic care management, transitional care, and other add-on elements are checked.

Stage Four: Code Generation. Only after Stage Three is complete does MDMai generate the final recommended billing codes. Because the upstream refinement has already happened, the codes are generated against refined clinical elements rather than against whatever the note happened to say in its raw form. The output is a structured report that lists each recommended code, the specific clinical language in the note that supports it, the CMS or payer rule that governs it, and any ambiguity flags where the documentation is close to the line.

Stage Five: The Biller Review. The MDMai report goes to the biller or practice manager. The biller decides what goes on the superbill. This is the architectural core of the system. MDMai does not write the superbill, does not submit claims, and does not remove the human from the billing decision. The report is a tool. The human has the last word on every claim that leaves the practice.

- **Why Human-in-the-Loop Is Deliberate, Not Incidental**

The decision to position MDMai between the signed encounter note and the superbill, rather than automating the superbill itself, is not just a design preference. It is a compliance architecture.

OIG compliance guidance issued in November 2023, and the Medicare Advantage industry compliance program guidance released in early 2026, have made explicit that AI-generated coding suggestions must receive meaningful human review before claim submission, that "automation bias" (the tendency for reviewers to rubber-stamp AI output) is itself a compliance risk, and that auditable justification linking every submitted code to specific clinical documentation is an expectation. The Department of Justice has identified AI-enabled billing workflows as an emerging False Claims Act enforcement priority. Healthcare recoveries accounted for approximately $5.7 billion of the roughly $6.8 billion in FY 2025 FCA settlements.

MDMai is built to fit cleanly inside that compliance frame. Every recommended code is traceable to specific passages of the note. The biller's review is the required human step. If the biller accepts the recommendation, the clinical documentation supporting it is already attached. If the biller rejects it, the reasoning sits in the audit trail. A practice using MDMai is positioned to demonstrate exactly what auditors and enforcement reviewers are looking for. A fully autonomous AI biller would not just be clinically risky. It would invite precisely the enforcement exposure the current environment is clearly moving toward.

- **The Seven Functions MDMai Combines**

Put the stages above together and MDMai performs, in a single integrated pass, seven functions that to our knowledge no other single system in healthcare combines:

Ingests encounter notes from any EHR. EHR-agnostic by design, with no requirement for API access, IT cooperation, or vendor approval.

Selects the most defensible allowable billing path (time-based or MDM-based) for each eligible E/M encounter under the 2021 and current guidelines.

Reconstructs clinician time on the date of service and identifies prolonged-services and add-on code opportunities that are routinely missed.

Refines every billable element (diagnoses, procedures, HCPCS, critical care criteria, chronic care management) against the clinical data in the note before any billing codes are generated.

Generates a structured, auditable report linking every recommended code to the specific documentation that supports it, giving the biller everything a compliance review would ask for.

Produces the annotated training dataset that ENCOUNTERai requires, paired note by note with MDMai's analysis.

Generates a payer-specific first-pass appeal automatically the moment a claim is denied or downcoded. The biller uploads the denial, MDMai maps the denied elements to the documentation already in the note and to the specific CMS or payer rule the denial appears to violate, and produces a fully assembled appeal letter in seconds. This inverts the appeal economics described in Part 1: instead of 30 to 60 minutes of biller labor per denied claim, the first-pass appeal is generated in seconds, which makes it economically rational for billers and practices to contest every improper denial rather than letting the intimidation loop force them to stop coding accurately.

To our knowledge, no EHR performs this kind of upstream MDM and element-specificity analysis out of the box, no AI ambient scribe analyzes existing notes for billing optimization, no CDI tool handles outpatient E/M MDM analysis systematically across payer rule sets, and no revenue cycle platform processes notes before coding occurs. The gap is as much structural as it is technical. We are not aware of another organization that has built a system designed to simultaneously serve as a billing optimization tool, a documentation quality analyzer, an AI training data pipeline, a compliance audit trail, and an automated appeal generator, and to do all of it while leaving the human biller firmly in charge of the superbill.

The Accidental Single Source of Truth

There is a second, quieter consequence of MDMai's architecture that was not a primary design goal but becomes one of the platform's most important properties in practice. Because MDMai is EHR-agnostic, because it ingests every signed encounter note a provider produces regardless of which system the note was signed in, and because every analysis output is written back to AHI's Healthcare Analytics Intelligence (HAL) platform under appropriate consent and privacy controls, the provider who uses MDMai across every site they practice is, without intending to, assembling the unified longitudinal record the government declined to mandate in 2009.

In the four-EHR scenario described in Chapter 2, the same patient's clinic note, two separate Cerner hospital notes, and the FQHC note would each pass through MDMai. The analysis output for each encounter, along with the underlying note, would sit in one place. For the first time in that patient's relationship with that provider group, there is a single store of record the treating clinician can consult at the point of care, keyed to the patient rather than to the facility that owns the EHR. MDMai was built to make notes billable. The side effect is that it also makes them findable.

This is not a claim that MDMai replaces an EHR, a Health Information Exchange, or TEFCA. It does not. The providers still chart in their hospital's Cerner or their clinic's system of record, and the legal medical record remains where the institution keeps it. What MDMai produces is a working, patient-centered, cross-facility view of every encounter the provider has signed, which is what is clinically useful at the bedside and what has been structurally missing from

American medicine since digitization began. In a market where interoperability has remained a promise for seventeen years, a unified view that emerges as a byproduct of solving the billing problem is an advantage worth being explicit about.

ENCOUNTERai: The Predictive Documentation Engine

ENCOUNTERai represents a fundamentally different approach to clinical documentation. Rather than transcribing what a physician says (ambient scribes) or templating what they click (EHRs), ENCOUNTERai is designed to generate complete encounter notes predictively from minimal structured input (diagnosis codes, key findings, treatment decisions) in each provider's clinical voice, optimized for billing from the moment of creation.

The underlying insight: research suggests that a substantial share of clinical note text is not originally composed at each visit. Wang et al. (JAMA Internal Medicine, 2017) found only 18 percent of progress note text was manually entered; 46 percent was copied and 36 percent imported from structured data. Rule et al. (JAMA Network Open, 2021) found 58.8 percent median note redundancy in an analysis of 2.7 million outpatient notes. Rotenstein et al. (JGIM, 2023) analyzed EHR usage patterns across large physician cohorts (sample sizes vary across the paper's analyses, including an 89,718–physician subset used for some specialty-specific analyses and roughly 215,207 physicians in the broader data pull) and documented strong provider-level patterns in documentation behavior. This predictability is precisely the condition under which a provider-specific AI model can be most useful.

The design goal is to approach approximately one minute of provider review-and-sign time per encounter, compared with the 16 to 36 minutes currently consumed by full documentation. This one-minute figure is a design target based on AHI's internal modeling, not a measured outcome in a clinical trial, and will be validated as the platform is deployed. The physician's review time is intended to be spent on clinical accuracy verification rather than billing compliance work.

The Training Pipeline

The MDMai-to-ENCOUNTERai pipeline operates in four stages.

Stage 1. Provider uses MDMai. Encounter notes are processed, providing immediate billing optimization and documentation feedback. The provider benefits from Day 1 through recovered revenue from undercoding and reduced denial risk.

Stage 2. MDMai generates annotated training data. Each processed note produces a paired dataset of raw note plus MDMai analysis (billing codes, MDM level, optimization recommendations, gap identification), stored in AHI's Healthcare Analytics Intelligence (HAL) platform under appropriate consent and privacy controls.

Stage 3. Once a provider has accumulated a sufficient corpus of annotated notes, ENCOUNTERai training begins. The accumulated dataset allows per-provider fine-tuning. The model is intended to learn the provider's documentation patterns, clinical reasoning, terminology, examination conventions, and the billing-to-documentation relationship. Published machine learning guidance (recent arXiv and Mayo Clinic Proceedings work) suggests that on the order of 1,000 to several thousand labeled samples can support reliable fine-tuning of clinical language models; AHI's current working threshold is in that range and is calibrated per specialty. These thresholds

reflect internal design targets rather than published outcome data, and will be refined as real-world evaluation proceeds. At typical outpatient volumes of 20 to 25 patients per day, providers can be expected to reach this threshold in roughly 60 to 90 working days.

Stage 4. ENCOUNTERai generates predictive documentation from minimal input, optimized for billing from inception. The provider reviews and signs.

This pipeline is why MDMai is offered free to physician practices. It is not philanthropy. It is the data acquisition mechanism that makes ENCOUNTERai possible. The value proposition is aligned. The provider benefits immediately, and the accumulated data creates an asset that would be difficult for competitors to replicate.

The Competitive Moat: Why This Pipeline Is Difficult to Replicate

EHR vendors have limited incentive to build a comparable pipeline. Their business model is platform lock-in, and a vendor-neutral analysis tool would threaten their position. AI scribe companies generate notes from voice but do not analyze existing notes or perform MDM analysis. Revenue cycle companies operate post-hoc. They process claims after submission, not clinical notes at creation. Academic institutions typically lack an EHR-agnostic processing engine deployable to any practice. The per-provider training threshold creates a time-based moat. Even a competitor who built a comparable engine tomorrow would need months of data per provider before their system could function.

The U.S. generates on the order of several billion clinical encounter notes annually. We believe AHI's MDMai platform is among the only EHR-agnostic mechanisms that can aggregate these notes with billing annotation at scale. The resulting database (encounter notes paired with MDM analysis, billing codes, and optimization data) is intended to become one of the more complete clinical documentation datasets available, and it grows with every provider who uses the free MDMai tool.

KEY POINT

Sources: AMA CPT 2021 E/M Revisions and FAQs; CMS Medicare Claims Processing Manual Transmittal R11842CP; CMS E/M Services Guide (September 2024); Arndt et al., Annals of Family Medicine (2017); AHIMA and ACDIS clinical documentation integrity guidance on sepsis and under-specificity coding; AAPC and Find-A-Code guidance on 99417, G2212, 99418, G0316, G0317 prolonged services; HHS OIG General Compliance Program Guidance (November 2023); HHS OIG Medicare Advantage Industry Compliance Program Guidance (2026); DOJ False Claims Act enforcement summaries (FY 2025); arXiv (2025); Databricks; Mayo Clinic Proceedings (2024); JMIR (2025); Wang et al. (2017); Rule et al. (2021); Rotenstein et al. (2023); AAPC (2024)

CHAPTER 5

THE 1,400-HOUR DIVIDEND

From 1,400 Hours to 100 Hours

A large AMA/Sinsky analysis of roughly 200,000 physicians across all specialties found that physicians spend, on average, several hours per 8-hour scheduled day on the EHR, with specialty-level variation. Widely reported figures place the all-specialty average in the range of 5 to 6 hours, with primary care and internal medicine specialties (including infectious disease) among the highest, in some analyses approaching or exceeding scheduled patient time. Exact specialty-level estimates (such as the often-cited 5.8, 7.3, and 8.4 hour figures) depend on the specific JGIM analysis cited and should be read alongside the underlying paper. Extrapolated across 240 clinical days, this translates to approximately 1,400 to 1,750 hours annually consumed by EHR documentation. Family physicians additionally log 86 minutes of after-hours "pajama time" each night, roughly 340 hours per year stolen from evenings and families.

METRIC
Current State
Target With ENCOUNTERai
Note creation per encounter
16 to 36 minutes
Design target: ~1 minute review and sign
Daily documentation
4 to 6 hours
Design target: ~25 to 30 minutes (25 pts)
Annual documentation time
1,100 to 1,600 hours
Design target: ~100 to 120 hours
Time returned
Design target: ~1,000 to 1,500 hours annually
After-hours pajama time
86 min/day (~340 hrs/yr)
Design target: near zero

Target figures above reflect AHI's internal modeling of ENCOUNTERai once trained for a given provider and are not outcomes from a completed clinical trial.

A time dividend of 1,000 to 1,500 hours would create meaningful physician-directed optionality. At 18 minutes of saved documentation time per encounter, recaptured hours could in principle support additional patient visits per year. As an illustrative projection only, a reallocation of 5 additional patient visits per day could translate to roughly $150,000 in additional annual revenue per physician at average Medicare reimbursement; actual outcomes will vary substantially with specialty, payer mix, local fee schedules, and physician capacity, and have not been measured in a controlled study. Across one million physicians, even 500 additional visits per physician equals 500 million additional patient encounters nationally, a substantial offset to the AAMC's projected shortage of up to 86,000 physicians by 2036, with a separate HRSA analysis projecting shortfalls exceeding 187,000 by 2037.

Roughly three-quarters of physicians, in recent surveys, would trade compensation for time, supporting the view that documentation burden is a binding constraint, not just compensation level.

The Medicaid Access Unlock

The economic barrier to seeing Medicaid patients is not low reimbursement alone. It is low reimbursement combined with high documentation time. A Medicaid 99213 visit pays roughly $55 to $66 nationally (reported as low as approximately $34 in some low-rate states such as Florida). When that visit requires 20 to 30 minutes of documentation time, the implicit hourly rate for documentation time alone is in the range of $110 to $198. With ENCOUNTERai, AHI's design target is for the same visit to require roughly one minute of review (again, a design target rather than a measured outcome). The economic friction that drives 26 percent of physicians to refuse new Medicaid patients is significantly reduced.

The downstream cascade: more physicians accept Medicaid patients, FQHC demand decreases (reducing the $32 billion parallel system cost), avoidable ER visits decrease ($32–plus billion annually), wait times decrease (currently around 31 days on average, up significantly since 2004 according to AMN Healthcare's 2025 physician wait time survey), and patient outcomes improve through earlier access to primary care.

The Documentation Quality Paradox: Less Time, Better Notes

Olson et al. (JAMA Network Open, 2025) demonstrated that when AI documentation tools partially fix the burden, burnout drops from 51.9 percent to 38.8 percent in 30 days, after-hours time drops approximately 0.90 hours per day, and note-related cognitive task load falls 2.64 points. Reduced burden did not reduce note quality.

ENCOUNTERai separates clinical quality from billing optimization. Today, physicians must simultaneously think about what is clinically relevant, what is legally protective, and what supports billing. This triple cognitive load can degrade all three. ENCOUNTERai is designed to handle the legal and billing dimensions automatically, allowing physicians to focus on clinical accuracy during their brief review. The goal is notes that are better clinically despite requiring much less physician time.

KEY POINT

Sources: Sinsky et al., JGIM (2024); Arndt et al. (2017); AMN Healthcare Physician Wait Time Survey (2025); Olson et al. (2025); MACPAC; AAMC; HRSA

CHAPTER 6

PAYER RETALIATION, PAY BY INTIMIDATION, AND THE MDMai DEFENSE

A natural objection arises. If MDMai reduces undercoding and practices suddenly bill at accurate levels, will insurers respond with more aggressive denials, audits, and downcoding? The evidence shows this concern is real but strategically manageable. Payers are already retaliating against accurate coding, making MDMai's defense architecture not optional but essential.

Part 1 of this white paper described the closed loop that the payer retaliation exploits: accurate coding triggers denials, appeals are individually unprofitable for billers working on commission (30 to 60 minutes of labor per appeal against a 7 to 10 percent commission on a recovered claim), and practices therefore come under pressure to stop coding accurately to preserve the biller's economics. We call that closed loop pay by intimidation. MDMai's defense architecture is designed to neutralize the loop by automating the appeal itself, which inverts the appeal economics that the payer is relying on.

The Cigna / Aetna Downcoding Offensive

Effective October 1, 2025, Cigna's Evaluation and Management Coding Accuracy Policy (R49) and Aetna's Claim and Code Review Program introduced automated review of Level 4 and 5 E/M visits, in which payment is adjusted downward when the submitted diagnosis code is deemed not to support the billed complexity level without a review of the underlying medical record. (Some trade press coverage in late 2025 reported that Cigna temporarily paused R49 amid provider backlash before reinstating it; readers should consult primary Cigna policy bulletins for the current status.) To receive the originally billed payment, providers must file a post-payment appeal with full medical records. E/M code 99214 is the most frequently billed E/M code in Medicare Part B, representing approximately 25 percent of all allowed claims. ASIPP concluded that "overcoding of CPT codes 99214 and 99215 is de minimis," suggesting that these policies target legitimate coding rather than fraud.

ASIPP characterized the policies as a "deny first, pay later" system that "presumes physicians are guilty of 'upcoding' until proven innocent." The strategic logic is transparent. Insurers expect that busy practices will not have the time or resources to appeal every improperly downcoded claim, resulting in net savings. ASIPP also noted that independent practices lack the armies of coders and administrators that hospital systems have, and often lack the capacity to engage the insurers' bureaucratic processes.

40-Plus Health Systems Have Called the Bluff

More than 40 health systems have dropped or threatened to drop Medicare Advantage contracts in 2024 and 2025, including Mayo Clinic, NewYork-Presbyterian, UNC Health, HealthPartners, and Baylor Scott & White. Providers often have more leverage than they realize. HealthPartners alleged UnitedHealthcare denial rates were "up to 10 times higher than other insurers." A KFF report found that MA enrollees had access to only about 48 percent of the physicians available to Traditional Medicare beneficiaries.

CMS has been widely reported to have never imposed intermediate sanctions on Medicare Advantage plans for network adequacy violations. MedPAC's June 2024 report noted that CMS has the authority to impose intermediate sanctions or civil monetary penalties for noncompliance with network adequacy standards, but it has never done so. Provider exits force payers either to renegotiate or to risk losing members when health systems grant patients special enrollment periods to switch plans.

Why Documentation Beats Algorithms

The evidence on claims appeal success rates is central to MDMai's defense:

METRIC	RATE	SOURCE
Internal appeal overturn rate (commercial)	~44 percent	KFF, 2023
Medicare Advantage denial overturn rate (appealed)	~57 percent	Health Affairs, 2025
Medical necessity denial overturn (with docs)	~70 percent	Industry analysis
MA prior auth denial overturn rate	~75 percent	HHS OIG, 2018
Cigna MA prior auth overturn (appealed)	~80 percent	House E&C Committee
Claims actually appealed (share of denials)	<1 percent	KFF, 2023

These figures come from different datasets and denial types and should be read individually rather than combined into a single range. They support a consistent conclusion: most of the denial system's economic advantage rests on low appeal rates, not on the strength of individual denials.

The system's most exploitable vulnerability is that fewer than 1 percent of denied claims are appealed. The payer strategy depends on this: deny at scale, count on providers not appealing, and retain revenue on the overwhelming majority of denials that go unchallenged. MDMai is designed to disrupt this equation. When documentation is airtight from creation and auto-generated appeal letters map specific MDM elements to CMS guidelines, AHI's design goal is to materially reduce the provider's cost of appealing; this is a design target rather than a measured outcome to date. If achieved, the payer's "deny first" strategy becomes economically less attractive at typical denial volumes.

The Auto-Appeal Generator: Neutralizing Pay by Intimidation

Part 1 of this white paper named the closed loop that keeps providers undercoding: payers deny accurate claims, billers lose money appealing at standard commission rates, and practices stop coding accurately to protect the biller's economics. Pay by intimidation. The auto-appeal generator inside MDMai is the technical response to that loop.

When a claim is denied or downcoded, the biller uploads the denial notice. MDMai already has the underlying encounter note in memory, along with the MDM analysis it produced when the claim was first coded. The system generates a payer-specific appeal letter in seconds, citing the exact documented elements that support the originally billed level, the specific CMS policy or payer medical policy language that contradicts the denial rationale, and the precedent-style argument appropriate to that denial reason. The biller or practice reviews, approves, and submits. What previously cost 30 to 60 minutes of labor per claim now costs a few minutes of review.

The economic math inverts. At 30 to 60 minutes of labor per claim, appeals were individually unprofitable for billers working on commission, which is why providers were pressured to stop coding accurately. With automated first-pass appeal generation, the labor cost per appeal drops sharply, which means appeals become rational for the biller and the practice to pursue. At that point, the payer's denial strategy stops being a cheap revenue-retention mechanism and starts being a cost center, because each denial now draws a well-documented appeal that the payer has to actually adjudicate. When the cost of denying accurately coded claims exceeds the expected recovery from providers who give up, the denial machine loses its economic rationale.

The auto-appeal generator is not a workaround for the payment architecture. It is a direct counter to the enforcement mechanism that the denial system uses to keep the payment architecture broken. It is specifically the reason MDMai is designed as both an offensive tool (coding accuracy) and a defensive tool (appeal automation) in a single platform.

Policy Lever: States That Raise the Wage Floor Must Also Raise Reimbursement

There is also a state-policy corollary. States that raise the minimum wage for healthcare workers (California's SB 525 is the clearest current example) create a cost-reimbursement scissor that accelerates the closure of independent practices unless they pair the wage increase with an equivalent update to Medicaid physician fee schedules and commercial parity protections. The scissor is not an argument against wage increases for healthcare workers. It is an argument that wage policy and reimbursement policy must move together, because the note is the only point where the combined cost side and revenue side of the practice meet. Technology alone cannot absorb a permanent scissor if the state raises cost while the payer holds reimbursement flat.

KEY POINT

Sources: Cigna R49 policy documentation; ASIPP (2025); Becker's Hospital Review (2025); KFF Health News (2025); MedPAC (2024); KFF (2023); HHS OIG (2018); House Energy & Commerce Committee oversight materials; California SB 525 (2023); U.S. Department of Labor state wage data; MGMA benchmarks on appeal labor cost

CHAPTER 7

THE UPWARD CASCADE

Part 1 traced the downward cascade: encounter note, documentation burden, burnout, declining reimbursement, access disparities, practice collapse, corporate consolidation, taxpayer workarounds, ER misuse, industrialized denial, and $5.3 trillion in total spending. A key insight is that the cascade is not linear. It is exponential. Each link can improve every other link when reversed.

Corporate Consolidation Reversal

Physicians overwhelmingly prefer independence. The Bain & Company 2024 Frontline of Healthcare Survey reported that physicians at hospital-led organizations expressed markedly lower satisfaction than those at physician-owned practices, with net promoter scores well below those of independents. Small independent practices show burnout rates of roughly 13.5 percent in one 2018 study, versus roughly 54 percent in the national physician average at the time, a fourfold difference.

The AMA found the top reasons physicians sold practices were the need to negotiate better payment rates (70.8 percent), the need to manage regulatory and administrative requirements (64.9 percent), and the need to access costly resources (63.6 percent). MDMai directly addresses the first (revenue optimization), and ENCOUNTERai is designed to reduce the documentation driver of the second. When independent practice becomes economically viable again, the consolidation pressure can reverse.

The Physician Pipeline Recovery

42.9 percent of current physicians would not choose medicine again, according to Physicians Foundation data. Nearly two-thirds would not recommend healthcare as a profession. Among physicians under 45, self-employment fell from 44.3 percent in 2012 to 31.7 percent in 2022. Only 22 percent of medical residents anticipate owning a practice stake. The AAMC projects a shortage of up to 86,000 physicians by 2036, with a separate HRSA analysis projecting shortfalls exceeding 187,000 by 2037.

The pipeline collapse is not primarily driven by compensation. It is driven by the working conditions documentation burden creates. When aspiring physicians see more than $200,000 in medical school debt, declining real pay, lost autonomy, liability risk, and over 1,000 hours per year of data entry, the calculation is rational. Many choose not to

enter medicine. When ENCOUNTERai returns those hours, the career proposition changes. The same physicians who currently regret their career choice would, in a system where they practiced actual medicine for 6 to 7 hours per day and went home to their families in the evening, be more likely to recommend the profession again.

The NP and PA workforce is expanding rapidly, and NPs and PAs now deliver a substantial share of U.S. primary care in some analyses. Much of that growth reflects genuine value and patient access, particularly in rural and underserved areas, and is not the subject of this paper. AHI's argument is narrower: to the extent physician substitution is being driven by an economic collapse of physician practice under documentation burden, reversing that burden changes the substitution math. Peer-reviewed research on private equity hospital acquisitions, including work led by Harvard Medical School investigators and published in high-impact medical journals, has found increases in in-hospital adverse events and patient-safety signals (including excess falls and infections) after acquisition, alongside reductions in staffing.

The $40+ Billion Workaround Unwind

The government currently spends tens of billions annually on workarounds for the access crisis: roughly $32 billion on the FQHC system (1,359 awardees, more than 17,000 sites, 32.4 million patients, and 139 million visits) plus approximately $9 billion in state supplemental physician payments (MACPAC, 2021 analysis). FQHCs receive $150 to $300 or more per Medicaid visit through the Prospective Payment System, while private practices receive roughly $55 to $66 for a comparable visit.

When documentation burden is significantly reduced and Medicaid patients become economically viable for private practices, the demand for the FQHC parallel system can diminish. These workaround budgets are not eliminated. FQHCs serve millions of uninsured patients and provide wraparound services. But the economic logic of continually expanding a parallel system to compensate for broken reimbursement weakens when the mainstream provider network can sustain Medicaid panels again.

The Single Record: Lives and Limbs Saved

The billing case for fixing the encounter note is large. The clinical case for a single, unified patient record, assembled as a byproduct of that fix, is larger. Picture a near-term future in which a patient's full longitudinal history, every encounter note across every EHR and every facility they have ever been seen at, is available at the point of care to the clinician treating them right now. That future is the direct consequence of deploying an EHR-agnostic platform like MDMai broadly. It is not speculative. It is what the platform already produces for every provider who uses it. Scaled, it becomes the repository that HITECH declined to require in 2009.

Consider what that single record changes. The emergency physician who receives an unresponsive patient by ambulance sees, within seconds, the full medication list, the active problem list, the allergy history, the recent lab trend, the most recent imaging, the primary care note from two weeks ago, the cardiology consult from three months ago, and the discharge summary from the admission last year at the hospital across town. The medication reconciliation error rate, measured in the current environment at above 60 percent on admission, collapses toward

something that resembles accuracy. The 41 percent allergy-list discrepancy rate collapses toward zero. The sepsis clock starts on presentation because the antibiotic history and the recent culture data are already visible, rather than on the fortieth minute after the ED team stops trying to reach the outpatient practice by phone.

Consider the transfer. The patient is transferred from Hospital A to Hospital B because Hospital B has the specialty capacity. Under today's architecture, the receiving team works from a fax packet that is, in the published literature, incomplete in roughly 40 percent of cases and missing material clinical information in roughly one-third. In the single-record future, the receiving team reads the same chart the sending team was charting in, with no seam, no rekeying, and no reliance on a patient too sick to narrate their own history. The 17 percent in-hospital mortality figure for transferred ICU patients is not a fixed constant of medicine. It is a function of how much the receiving team does not know. Close that information gap and a measurable share of those deaths become preventable.

Consider the limbs. Vascular complications, diabetic foot, necrotizing infections, compartment syndromes, and ischemic limbs all share a feature: time is tissue, and the prior history (vascular study from last quarter, prior amputation, prior antibiotic exposure, anticoagulation status, renal function) determines the speed and the correctness of the intervention. A clinician who has that history in front of them within seconds makes a faster, sharper call than a clinician reconstructing it from scratch at 2 a.m.

Consider the medication-error death toll, which is a substantial component of the Makary and Daniel estimate. A large share of medication errors arise from not knowing what the patient is already on, what they have reacted to before, what has been discontinued and why, what dose adjustments the nephrologist made last month, and what the last ED visit prescribed. None of that is mysterious. All of it is already written down somewhere. It simply is not, today, findable from where the next clinician stands. A single repository changes that.

Consider the chronic disease patient. The diabetic who sees a primary care physician, an endocrinologist, a podiatrist, a cardiologist, and an ophthalmologist experiences five separate partial views of their condition under the current architecture. Each specialist sees a fraction. No single clinician in the patient's life has ever seen the whole picture. A unified record gives each of them the same view, which makes every decision downstream more coherent. Preventable complications (the stroke that followed missed anticoagulation, the kidney failure that followed an unreconciled NSAID, the amputation that followed a podiatry note the hospitalist never saw) become preventable.

Consider the litigation footprint. CRICO and Candello's communication-failure share of malpractice claims (30 percent, rising toward 40 percent in more recent analyses) largely describes claims that would not exist if the information had reached the clinician in time. A real single record does not eliminate malpractice. It meaningfully shrinks the surface area of the specific category of claims that is currently the largest and fastest-growing.

Consider public-health surveillance. Unified records, queried with appropriate consent and privacy controls at population scale, would let epidemiologists see outbreaks, adverse drug events, and care-gap patterns in days rather than quarters. The COVID-era delay in linking vaccination status to outcome data across fragmented systems would not repeat in the same form.

The benefits run in one direction. The costs are not zero (data governance, consent architecture, patient matching at scale, cybersecurity, access control, and auditability all require real engineering and real policy), and nothing in this section claims those problems are solved on the back of an envelope. But the core observation holds. A platform that has to process every signed encounter note in order to make it billable is, structurally, already collecting the dataset that solves fragmentation as a side effect. The question is not whether the unified record is possible. It is whether the country lets a platform with the right architecture produce it, or whether the same failure mode that HITECH left unaddressed in 2009 is allowed to persist for another seventeen years.

This is the second reason AHI's technology is worth deploying broadly. The first reason is that it fixes the billing and burden problem at the practice level. The second reason is that it produces, as a byproduct, the single source of truth American medicine has needed from the moment the paper chart was retired. Lives and limbs are saved in the second reason. The first reason pays for the second.

KEY POINT

Sources: Bain & Company (2024); AMA (2024); Physician Advocacy Institute (2024); AAMC (2024); HRSA (2024); MACPAC (2025); KFF (2025); Physicians Foundation (2024); Joint Commission on handoff communication; AHRQ PSNet; Makary and Daniel, BMJ (2016); Critical Care Medicine interhospital-transfer study (2016); medication reconciliation literature; CRICO Strategies (2015) and Candello (2024-2025)

CHAPTER 8

THE POLICY AGENDA, WHAT THE TAXPAYER SHOULD DEMAND IN EXCHANGE FOR $2.4 TRILLION

The federal government will spend roughly $2.4 trillion on healthcare in 2024 and, on current CBO projections, roughly $4.3 trillion by 2033. That is taxpayer money. The taxpayer has been the silent underwriter of the private-public symbiosis described earlier: private insurers profit from the healthy commercial population; the federal and state governments cover the populations the private system systematically designs against; and every widening of the access gap, every post-acquisition price increase, every Medicare Advantage overcoding case, and every avoidable emergency department visit lands somewhere on the federal balance sheet.

AHI's technology strategy is deployment-first, for the reasons Chapter 1 set out: the encounter note is the one intervention point where technology can change practice economics without requiring the private-public split to be legislatively renegotiated. That strategy does not absolve the federal and state governments of their responsibility to reform what they alone can reform. The following is the list of specific asks the taxpayer should be making in exchange for the $2.4 trillion they currently provide. Each is grounded in evidence assembled in Part 1 or in publicly available oversight materials. Each is one the industry has lobbied against or delayed. And each, if enacted alongside broad deployment of tools that fix the encounter note, would compound in effect.

Ask One: Medicare Advantage Coding-Intensity Enforcement and Overpayment Recovery

MedPAC and HHS OIG have documented for years that Medicare Advantage risk scores run systematically higher than traditional Medicare for comparable populations because of coding intensity rather than actual morbidity. Estimated overpayments attributable to coding intensity ran approximately $23 billion in 2023 and are projected to exceed $40 billion by 2025. CMS applies a coding intensity deflation adjustment of approximately 5.9 percent, which is smaller than the documented gap. The taxpayer should demand enforcement action commensurate with the scale of the overpayment, meaningful recoupment of unsupported diagnoses (OIG audits have found substantial shares of audited MA diagnosis codes lack supporting documentation), and auditing protocols applied at the per-plan level rather than the aggregate level so that individual plans cannot hide behind industry averages.

Ask Two: Medicaid Physician Fee Parity With Medicare

MACPAC's Medicaid-to-Medicare physician fee index has remained in the 0.66 to 0.72 range for two decades, reported at approximately 0.71 in 2024. The 2013 to 2014 ACA Medicaid primary care fee bump briefly raised Medicaid primary care rates to Medicare parity and produced measurable increases in appointment availability before it was allowed to expire. The policy demonstrated that the fix worked. Permanent Medicaid physician fee parity with Medicare would eliminate the single largest structural driver of the access gap. The federal cost would be meaningful but would be offset in part by reduced federal outlays on the FQHC workaround (roughly $32 billion annually), state supplemental physician payments (roughly $9 billion annually per MACPAC's 2021 analysis), and Medicaid emergency department utilization (Medicaid patients use EDs at approximately 4.2 times the privately insured rate). No major federal analysis has compared the total cost of Medicaid parity against the combined cost of the current workaround system.

Ask Three: Prior-Authorization Reform With Teeth

Prior authorization is one of the most concrete documented transfers of administrative cost from payers to providers. CAQH estimated $1.3 billion in annual administrative cost, a substantial share borne by providers. The AMA's 2024 prior authorization physician survey documented widespread harm to patient care. HHS OIG's 2022 audit of Medicare Advantage prior authorization found substantial inappropriate denials. Narrow prior-authorization reforms have been proposed and advanced in multiple Congresses; what has not been enacted is a statutory framework with teeth. The taxpayer should demand: mandatory electronic prior authorization with statutory turnaround deadlines; transparent public reporting of prior-authorization denial and appeal success rates at the per-payer and per-service-category level; and civil penalties for payers whose denial patterns are overturned on appeal at rates above a statutory threshold.

Ask Four: Denial and Appeal Transparency

KFF's 2025 analysis of ACA marketplace claims denial found denial rates averaging roughly 19 percent of in-network claims, with appeal rates under one percent. HHS OIG's work on Medicare Advantage denials found substantial rates of inappropriate denial. ProPublica's reporting documented automated payer review systems issuing denials in seconds without clinical review. The denial economics described as pay by intimidation in Part 1 can only persist where the data on denial and appeal success rates are opaque. The taxpayer should demand that every payer operating in any federally subsidized market (including ACA exchange, Medicare Advantage, and Medicaid managed care) publish at the quarterly level: claim volume by service category; denial rate by service category and reason code; appeal filing rate; appeal success rate; and average time-to-adjudication. Denial at scale is not compatible with transparency at scale.

Ask Five: Narrow-Network Adverse Selection Oversight

Between 2014 and 2024, the share of ACA Exchange plans structured as narrow-network HMOs or EPOs rose from 42 percent to 79 percent, while broader-network PPOs fell from 58 percent to 21 percent. Commonwealth Fund research in 2025 documented that the lowest-premium marketplace plans, which lower-income households are most likely to select, also carry deductibles that make most routine care effectively unaffordable. These are design choices. They are not illegal under the ACA, but they are consistent with risk-selection behavior. The taxpayer, who subsidizes roughly 85 percent of marketplace premiums through advance premium tax credits, should demand that HHS and state insurance regulators publish adverse-selection metrics by plan, specifically including the ratio of enrolled chronic-disease-burden risk to the geographic market baseline. Plans that persistently underweight sicker populations should face rate-review consequences.

Ask Six: Public-Program Reimbursement Must Move With State-Level Cost Increases

California SB 525, as Part 1 described, raised the healthcare minimum wage beginning June 2024. Other states are considering similar legislation. Wage increases for healthcare workers are appropriate policy. What is not appropriate is enacting cost-side legislation without adjusting the reimbursement-side floor that public programs set. Every state that raises the healthcare labor cost floor should be required to submit, as part of its Medicaid State Plan Amendment process, a corresponding reimbursement adjustment for Medicaid physician fee schedules within the next rate-setting cycle. Otherwise the state-level wage floor compounds the closure rate of independent practices that serve Medicaid patients and accelerates the corporate consolidation the state's own residents will pay for downstream.

Ask Seven: Federal Workaround Cost Accounting

The federal government has never published a unified accounting of the total cost of the Medicaid-access workaround infrastructure. FQHC grant funding, Medicaid PPS payments to FQHCs, Disproportionate Share Hospital payments, uncompensated emergency care subsidies, Medicaid managed care directed payments tied to physician services, and state supplemental payment programs are accounted for separately. Aggregated, they represent tens of billions of dollars annually, built explicitly because the underlying Medicaid fee schedule does not support private-practice economics. The taxpayer should demand a single CMS- or CBO-produced annual report aggregating these costs and comparing them with the counterfactual cost of raising the underlying Medicaid fee schedule to Medicare parity. This comparison has never been done at the federal level. It should be.

The Deployment-First Strategy Does Not Replace These Asks

AHI's technology strategy, as Chapter 4 described, is to change the economics at the practice level without waiting for any of the above to pass. That is the correct strategy given the lobbying equilibrium described in Chapter 1. But deployment-first and policy-reform-next are complementary, not substitutes. The more independent practices deploy tools that fix the note, the more data exists to support each of the policy asks in this chapter. The more Medicaid panels reopen, the more visible the FQHC workaround cost becomes. The more auto-generated appeals are filed against payer denials, the more denial and appeal data the taxpayer can demand. The deployment creates the political ground truth that the policy reform depends on.

The taxpayer should not wait for Washington to fix the encounter note. The taxpayer should insist on the seven asks above, and in the meantime, support deployment of tools that work whether or not Washington ever acts.

KEY POINT

Sources: CBO June 2024 baseline budget projections; CMS National Health Expenditure Accounts (2024); MedPAC and HHS OIG on Medicare Advantage coding intensity and overpayments; MACPAC (2021, 2023, 2024, 2025); Polsky et al., NEJM (2015) on the Medicaid fee bump; AMA 2024 prior authorization survey; CAQH administrative cost report; KFF 2025 claims denials analysis; HHS OIG 2018, 2022 on Medicare Advantage denials; Commonwealth Fund 2025 marketplace affordability; Oliver Wyman 2024 on ACA narrow networks; California SB 525 (2023)

§ 11

The Complete Solution: Fix The Note, Fix The System

The evidence assembled across both parts of this white paper leads to a single conclusion:

> KEY POINT
>
> The clinical encounter note is simultaneously the root cause of healthcare's $5.3 trillion dysfunction and an intervention point where significant parts of the cascade can be reversed, because it is one of the few points where technology can operate without requiring legislative approval, regulatory reform, EHR vendor cooperation, or insurer consent.

The logic is straightforward:

CURRENT DOWNWARD CASCADE	REVERSED BY AHI TECHNOLOGY
Documentation consumes ~1,400 hrs/yr	ENCOUNTERai targets ~100 hrs/yr
Undercoding leaves 5 to 11 percent revenue loss	MDMai reduces undercoding at note creation
Burnout at roughly 49 to 63 percent	Olson et al. showed a 13-pt drop with a partial fix
26 percent refuse new Medicaid patients	Documentation barrier reduced
Tens of billions in FQHC and supplemental workarounds	Private practices absorb more Medicaid demand
$32B-plus in avoidable ER visits	Primary care access expands
~77.6 percent physicians corporate-employed	Independent practice becomes more viable
19 percent of claims denied, under 1 percent appealed	MDMai supports appeals with airtight documentation
42.9 percent would not choose medicine again	Career proposition improves

The healthcare industry spent approximately $744 million in 2024 lobbying on federal policy. The purpose of that spend, as Part 1 of this paper documented and Chapter 1 of this one reinforced, is to preserve the specific split between what the private system covers profitably and what the public system absorbs at a loss. The encounter note

is what sits underneath that split. Every Medicaid patient whose encounter note cannot generate a margin at $34 per visit is a patient the private practice must either subsidize, turn away, or sell its independence to survive serving. Every one of those choices pushes the cost further up the chain: to the state, which builds out FQHC capacity and supplemental physician payment programs; to the federal government, which funds the FQHC system, Medicaid, Medicare, Medicare Advantage, DSH, and EMTALA-mandated uncompensated emergency care; and eventually to the taxpayer, who underwrites all of it. The private payer stays profitable because the public sector keeps catching what the private sector drops. That is what the lobbying spend protects.

The deployment-first strategy does not ask Washington's permission to begin changing this arithmetic. Insurers and their trade associations cannot lobby to prevent a physician from using a tool that saves them time and recovers revenue. They cannot lobby to prevent the encounter note from being written correctly. They cannot lobby to prevent a biller from filing an auto-generated appeal that is more rigorous than the cursory denial it answers. They cannot lobby to prevent burnout from dropping meaningfully in 30 days when documentation burden is reduced. They cannot lobby to prevent an independent practice from reopening its Medicaid panel once the math finally works.

Every practice that deploys a tool that fixes the note weakens the economic foundation of the private-public split. Every Medicaid panel that reopens in a private practice is one more patient the FQHC system does not have to absorb. Every auto-generated appeal filed against an automated denial is one more drag on the denial machine's unit economics. Every avoidable emergency department visit that does not happen because the patient could get an appointment within a week is one less dollar of federal uncompensated care expense. The snowball that Part 1 described runs in reverse the moment the encounter note stops being the primary constraint.

The path around a lobbying wall is through the encounter note itself. That is what we built. Chapter 8 of this paper describes the policy asks the taxpayer should make in parallel, because deployment-first and policy-reform-next compound when they run together. But the policy asks are not a prerequisite. The technology works whether or not Washington ever acts. It works at the practice level, in the encounter, on the claim, on the appeal, one note at a time, and every note that is written correctly is one more link in the cascade that bends.

> **KEY POINT**
>
> We didn't invent this evidence. JAMA published it. The AMA documented it. CMS data confirms it. The GAO reported it. We connected the dots that nobody else connected, and then we built the technology to fix it.

ABOUT

About Artificial Healthcare Intelligence

Artificial Healthcare Intelligence (AHI) is a 501(c)(3) nonprofit healthcare technology organization headquartered in the United States. Founded by John Leoniak, a software engineer with over 30 years of experience and more than a decade in healthcare operations, AHI was born from watching the crisis happen firsthand: his wife, Dr. Jennifer Leoniak, a board-certified infectious disease specialist whose independent infectious disease practice serves a referral-area population on the order of several hundred thousand residents, was spending more time on documentation than with patients and family.

During the COVID-19 pandemic, the practice's daily patient encounters expanded from 10 to 15 per day to 100 to 120 per day for three consecutive years, generating more than 100,000 encounter notes involving most medical specialties. This experience provided unique insight into both the documentation burden destroying physician quality of life and the extraordinary clinical value locked within encounter notes that no one was systematically capturing.

AHI's product suite is at varying stages of maturity. MDMai (medical documentation and billing optimization) and HAL (Healthcare Analytics Intelligence) are the flagship platforms under active development and early deployment. ENCOUNTERai (predictive documentation), HALi (patient-facing Healthcare Intelligence), CODEai, and REFERRALai are in earlier stages of development or planning, with capabilities described in this paper reflecting AHI's design goals rather than generally available commercial products. Readers evaluating the organization should treat all product claims in this paper as forward-looking unless otherwise noted. Dr. Jennifer Leoniak serves as Chief Medical Officer, providing the clinical validation essential for healthcare intelligence systems that demand high accuracy.

Contact

For media inquiries, partnership discussions, or additional information:

Artificial Healthcare Intelligence (AHI)

SOURCES

References

The following sources were consulted in the preparation of this white paper. All statistics, findings, and claims are traceable to the original publications cited below.

1 AAPC. (2024). AI Scribes: Coding Naive Analysis. October 2024.

2 ACDIS (Association of Clinical Documentation Integrity Specialists). Clinical documentation integrity guidance on sepsis specificity, pneumonia organism specificity, diabetes complication coding, and heart failure staging.

3 AHIMA (American Health Information Management Association). Sepsis Under the ICD-10-CM Microscope; clinical documentation improvement guidance on under-specification of diagnoses.

4 AMA. (2021). CPT Evaluation and Management (E/M) Office or Other Outpatient Services Code and Guideline Changes, effective January 1, 2021; E/M Revisions FAQs.

5 AMA. (2023). Hold Health Plans Accountable for Physician Network Inadequacies. November 2023.

6 AMA. (2024). Physician Practice Benchmark Survey.

7 AMA. (2025). Competition in Health Insurance: A Comprehensive Study of U.S. Markets.

8 AMA. CPT Professional Edition. American Medical Association, current edition.

9 AMN Healthcare. (2025). Physician Wait Time Survey.

10 Arndt, B.G. et al. (2017). Tethered to the EHR: Primary Care Physician Workload Assessment. Annals of Family Medicine, 15(5), 419–426.

11 AHRQ Patient Safety Network (PSNet). Transitions of Care primer and related analyses on communication breakdowns as contributors to preventable adverse events.

12 Asgari et al. (2025). Hallucination Rates in AI-Generated Clinical Notes. npj Digital Medicine.

13 ASIPP. (2025). Aetna and Cigna Announce Automatic Downcoding Policies. October 2025.

14 Bain & Company. (2024). Frontline of Healthcare Survey.

15 Becker's Hospital Review. (2025). 40 Health Systems Dropping Medicare Advantage Plans. November 2025.

16 Bipartisan Policy Center. Research on U.S. health care provider consolidation.

17 CMS. (2021). Evaluation and Management Services Guidelines.

18 CMS. (2024). Evaluation and Management Services Guide. Centers for Medicare & Medicaid Services, September 2024.

19 CMS. Medicare Claims Processing Manual Transmittal R11842CP. Prolonged services and E/M guidelines.

20 Commonwealth Fund. (2024). Mirror, Mirror 2024: A Portrait of the Failing U.S. Health System.

21 Critical Care Medicine. (2016). Adverse Events, Mortality, and Documentation Completeness in Interhospital ICU Transfers.

22 CRICO Strategies. (2015). Malpractice Risks in Communication Failures: Annual Benchmarking Report.

23 Candello. (2024–2025). Communication Failures in Malpractice Claims: Updated Benchmarking Analyses. The Risk Management Foundation of the Harvard Medical Institutions.

24 DOJ. (2025). False Claims Act Enforcement Statistics, Fiscal Year 2025. U.S. Department of Justice, Civil Division.

25 EPIC for America. (2024). Health Insurance Revenues Tripled in a Decade.

26 FTI Consulting. Research notes on Medicare Advantage payer-provider disputes.

27 Government Accountability Office (GAO). VA Electronic Health Record Modernization program reports.

28 Health Affairs. (2023). The Revolving Door in Health Care Regulation.

29 HealthLeaders Media. Reporting on Cigna and Aetna downcoding policies (2025).

30 HHS OIG. (2018). Medicare Advantage Organization Prior Authorization Denials.

31 HHS OIG. (2022). Medicare Advantage Inappropriate Claim Denials.

32 HHS OIG. (2023). General Compliance Program Guidance. Office of Inspector General, U.S. Department of Health and Human Services, November 2023.

33 HHS OIG. (2026). Medicare Advantage Industry Compliance Program Guidance. Office of Inspector General, U.S. Department of Health and Human Services.

34 House Committee on Energy and Commerce. Oversight findings on Medicare Advantage prior authorization overturn rates.

35 HRSA. (2024). Uniform Data System National Report.

36 Joint Commission. Sentinel Event and Handoff Communication analyses attributing a substantial majority of serious medical errors to communication failures during transitions of care.

37 Kaiser Permanente. (2025). AI Ambient Scribe Deployment Study. NEJM Catalyst.

38 KFF. (2023). Claims Denials and Appeals in ACA Marketplace Plans.

39 KFF Health News. (2025). Medicare Advantage Network Adequacy Enforcement. November 2025.

40 Lukac et al. (2025). Randomized Clinical Trial of AI Scribes. NEJM AI.

41 Makary, M. and Daniel, M. (2016). Medical Error, the Third Leading Cause of Death in the US. BMJ, 353.

42 MACPAC. (2024–2025). Supplemental Payment Analysis and Medicaid Fee Analysis.

43 MedPAC. (2024). Medicare Advantage Network Adequacy Report. June 2024.

44 MGMA. (2024). Practice Benchmark Data.

45 Office of the National Coordinator for Health IT (ONC). Data Brief 71 (May 2024): Electronic Exchange of Health Information Among U.S. Non-Federal Acute Care Hospitals.

46 Office of the National Coordinator for Health IT (ONC) and HHS. 21st Century Cures Act (2016) and Cures Act Final Rule on Information Blocking (2020).

47 Olson, A. et al. (2025). AI Ambient Documentation and Physician Burnout. JAMA Network Open.

48 OpenSecrets. (2024–2025). Federal Lobbying Data.

49 Palm et al. (2025). Hallucination Rates in LLM-Generated Ambient Notes. Frontiers in Artificial Intelligence.

50 Peterson Health Technology Institute. (2025). AI Scribe Financial Impact Assessment. March 2025.

51 Physicians Foundation. (2024). Physician Survey Report.

52 Physician Advocacy Institute. (2024). Physician Employment Trends Report.

53 Rotenstein et al. (2023). Documentation Strategies Among 215,207 Physicians. JGIM.

54 Rule et al. (2021). Changes in Outpatient Clinical Note Length. JAMA Network Open.

55 Sinsky, C. et al. (2024). EHR Time Across 200,081 Physicians. JGIM.

56 VA Office of Inspector General. Reports on Oracle / Cerner Electronic Health Record Modernization program.

57 Wang et al. (2017). Text Origin in Clinical Notes. JAMA Internal Medicine.