# State of AI in RCM 2024

Elion

# State of AI in RCM 2024

## A Message from our Head of Research

At Elion, we're passionate about helping healthcare leaders make better tech decisions through thorough, unbiased, and actionable research. This report reflects months of effort and hundreds of hours of work. We take our commitment to you, the tech buyer, seriously and appreciate you taking the time to read the report.

We believe 2024 will be a pivotal year for AI platforms in healthcare. While clinical use cases are intriguing—and we're eager to see how they progress—at the moment, AI is currently making the most impact in back-office or hybrid clinical/back-office scenarios. Workflow automation, autonomous coding, and ambient scribes are exciting applications with real-world potential.

From a market perspective, major players like Epic and Oracle Cerner haven't fully developed AI-powered solutions yet, creating a novel opportunity for a new kind of platform company. This, combined with the urgent business need for greater efficiency in the revenue cycle, may position RCM tech startups for potential unicorn status in the coming years.

Speaking of that urgent business need—I know I'm not telling you anything you don't know, but each year seems to bring new challenges. RCM in the U.S. is estimated to be a $156 billion market, yet the current way of doing business has its share of short-comings:

- 12% of all claims are denied.
- The cost to rework or appeal a denied claim is $25 for practices and $181 for hospitals.
- 65% of denied claims never get re-submitted.
- Hospitals often take 90 days or more to get paid by consumers and payers.

Given that health systems are losing between 1-5% of net patient revenue due to revenue cycle inefficiencies, AI's promises of lower costs and greater reimbursements are very appealing.

The question is: Do these promises live up to the hype? We can't tell you exactly which products to buy, as each organization's needs are unique. But we hope that after reading this report, you'll have a better understanding of how these technologies work, where they're delivering ROI today, and where the top vendors are heading in the next one to three years, empowering you to make better-informed decisions.

Sincerely,

P. B. Wingo

Patrick Wingo

# Report Scope and Definitions

Revenue cycle management includes all of the steps that ensure a healthcare provider organization is ultimately paid (see next page for a quick overview of the stages of RCM), and AI is applicable across nearly every step of this process.

But what does it mean to be an AI healthcare tech product? As technology evolves, so does the definition, but currently, "AI-enabled" in healthcare means one or more of the following:

- Using sophisticated data science for classifying or making decisions using numerical or categorical data.
- Using natural language processing (NLP) techniques for processing text or audio.
- Incorporating large language models (LLMs) or other deep learning techniques that require large amounts of training data.

In healthcare generally and RCM specifically, generative AI, NLP, and machine learning (ML) present major opportunities to improve revenue integrity by enabling the technology to:

- Parse unstructured data like clinical notes more efficiently than a human.
- Communicate with payers over the phone with minimal or no human oversight.
- Automate manual activities with little to no additional incremental cost.
- "Learn" from past payer exchanges to increase the probability of payment through better documentation, medical coding, and claims editing.
- Analyze large quantities of data to uncover potential opportunities for improvement and predict negative outcomes to prevent them from occurring.

Perhaps most compellingly, while clinical interactions rely on relationships and empathy, revenue cycle management does not require the same human touch. Revenue cycle activities also involve significantly less clinical, legal, or ethical risk. In short, it's a space ripe for AI-driven innovation.

In this report, we'll cover the revenue cycle from the perspective of an individual claim, examining the current workflow and identifying opportunities to apply AI. Our analysis will predominantly focus on the provider side of the equation (as opposed to the payer side) and will look at fee-for-service workflows, although there is some overlap with those involved in value-based care models.

## AI Terminology

- **Machine Learning (ML):**
  Field of AI focused on developing algorithms that can learn from existing data to perform tasks without explicit instructions.
- **Natural Language Processing (NLP):**
  A field of AI that uses machine learning to understand human language.
- **Large Language Models (LLMs):**
  Very large deep learning or machine learning models trained on specific (large) datasets to perform RAG (retrieval augmented generation) tasks. For example, medical LLMs could be trained on medical datasets, such as study findings or EHR data.
- **Generative AI (GenAI):**
  Type of AI that creates new media based on learning from existing data.

# The Life Cycle of a Claim

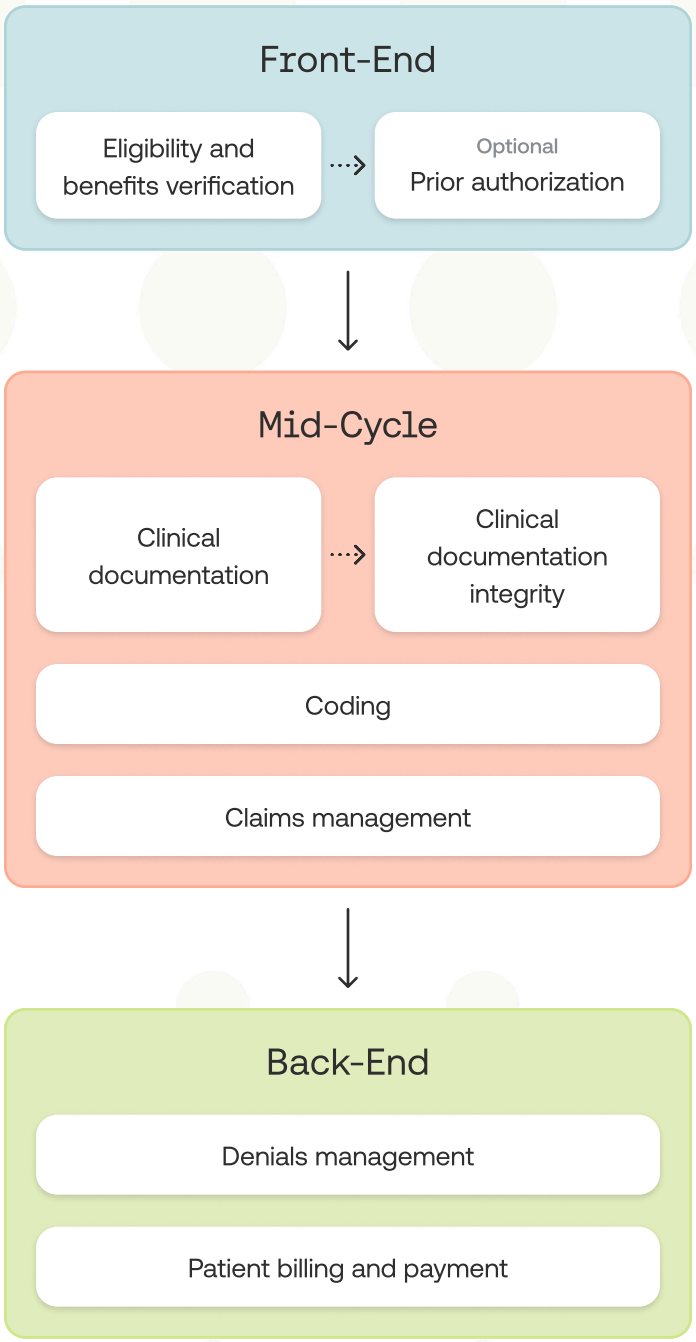

# AI in RCM Market Maps

Elion

## Front-End

### Revenue Cycle Automation

- Luminai
- Automat
- Element5
- Infinx
- Jorie AI
- Plenful
- Rotera Alyks
- SparkChange
- Thoughtful AI
- UiPath

### Payer-Facing AI Phone Calls

- Amperos Health
- Bland AI
- Infinitus
- Opkit
- SuperDial
- Health Harbor
- Majorboost
- Outbound AI
- Standard Practice
- Vapi
- Develop Health

### AI Prior Authorization for Providers

- AKASA
- Autonomize AI
- Infinitus
- Silna
- Synthpop
- Develop Health
- Ethermed
- Humata Health
- Latent
- Myndshft
- SamaCare
- Tandem
- Thoughtful AI
- Waystar

## Mid-Cycle

### AI Ambient Scribes

- Ambience Healthcare
- Nabla Copilot
- DeepScribe
- Nuance DAX
- Augmedix Go
- Suki Assistant
- Abridge
- MarianaAI

### AI Medical Coding

- Fathom
- Nym
- RapidClaims
- Revedy
- Aideo
- AKASA
- Arintra
- CodaMetrix
- CorroHealth
- Hank AI
- Maverick Medical AI
- Medicodio Codio
- Phare Health
- Semantic Health

### AI Clinical Documentation Integrity

- Ambience Healthcare
- MarianaAI
- SmarterDx
- 3M 360 Encompass
- 3M CDI Engage One
- Brellium
- Evidently CDI
- Iodine Aware CDI
- Optum
- Regard
- Semantic Health
- Accuity

### Claims Management

- Athenahealth
- Epic
- Oracle Cerner
- NextGen EHR
- Apex EDI
- Availity Essentials
- Change Healthcare
- Experian
- Trizetto
- TruBridge
- Veradigm
- Waystar
- Adonis Billing
- Apero
- Candid Health
- Docvocate
- Enter Health
- Finthrive
- Inovalon
- Optum
- pMD
- Trek Health RCM
- Optum

## Back-End

### AI Denials Management

- Adonis Intelligence
- Change
- Crosby Health
- Datavant
- Etyon
- Experian
- MDClarity Revfind
- MedMetrix
- Rivet Claims Resolution
- Rivet Payer Performance
- Sift
- SSI
- Waystar

### AI Patient Billing

- Cedar Pay
- Decoda Health
- Ecton
- Flywire
- PayGround
- PayZen
- Raxia
- Sift
- Vantage Health
- Clearbalance

Accurate as of **August 2024**. For the most up-to-date list of companies, visit Elion.

# Outline

# RCM Front-End



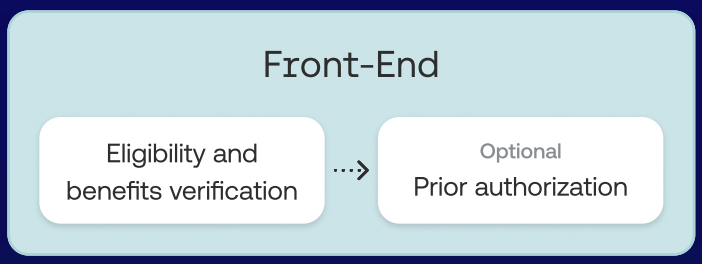


The front-end of the revenue cycle consists of all of the work that takes place before the patient is actually treated. However, this stage has profound implications throughout the revenue cycle, as eligibility and prior auth denials are often in the top three denial reasons for health systems' claims. Technically, some RCM leaders might also include patient access within RCM front-end, but for our purposes, we've organized this report to follow the progress of a single claim. And so, we begin with eligibility and benefits checking.

# Eligibility and Benefits Checking

In simple terms, *eligibility* refers to whether the patient is currently covered by an active health insurance plan, while *benefits* refers to details of the coverage for a particular treatment or service, including:

- Is the service covered?
- Are there medical necessity requirements?
- Does it require prior authorization?
- What level of cost-sharing is involved, including copays and coinsurance?
- Are there limits or exclusions?
- Where does the patient stand in relation to their deductible and out-of-pocket maximum?

## Current Workflow

There are two traditional workflows for checking eligibility and benefits: **software-based checks and manual checks via phone calls or payer portals** (the payer-specific online platforms where providers can verify patient eligibility, submit claims, and access other insurance-related information). Unfortunately, both options have their drawbacks.

EDI X12 software-based eligibility checks, typically performed through the EHR, are effective for routine verifications. However, they often fall short in handling complex scenarios or special cases.

Manual checks via phone calls or payer portals can support more complex eligibility situations, however they are highly labor-intensive and costly. Verifying the eligibility of a single patient can take several hours, requiring significant effort, specialized training, and substantial staff investment. Both methods are prone to human error, leading to inaccuracies and potential delays. Additionally, these methods lack the scalability and efficiency needed to handle high volumes of patient eligibility checks.

AI-enabled solutions are designed to address these shortcomings by automating the process of checking payer portals or making calls to payers, so that providers can get more accurate details on eligibility checks without the scalability challenges.

### **Eligibility and Benefits Check Workflows:** How and Why Each Might Be Used

#### Software X12 Eligibility Requests (EDI 270/271)

- **Standard Practice:**
  Health systems commonly send X12 270 requests to check patient eligibility electronically. This is part of the Electronic Data Interchange (EDI) standard that powers healthcare rails more broadly.
- **Automation:**
  X12 requests are often integrated into practice management or electronic health record (EHR) systems, allowing for automated, real-time eligibility checks.
- **Efficiency:**
  Electronic requests are quicker and more efficient, providing immediate responses and reducing administrative burden.
- **Batch Processing:**
  Health systems can process multiple eligibility checks in batch mode, which is useful for larger volumes.

### Manual Checks: Phone Calls and Payer Portals

- **Complex Cases:**
  Health systems may call payers directly or use payer portals for more complex eligibility verification needs that may not be fully addressed through electronic requests.
- **Clarification:**
  In situations where there is ambiguity or uncertainty in a response to an electronic check, a call can provide clarification.
- **System Downtime or Errors:**
  In cases of system downtime or technical issues with the EDI, health systems may resort to direct calls to ensure there are no delays in patient care.

### When One Method is Preferred Over the Other

- **Volume of Requests:**
  For high-volume eligibility checks, X12 is preferred due to its scalability and automation capabilities.
- **Type of Information Needed:**
  For basic eligibility and benefit information, X12 is usually sufficient. For more detailed, nuanced, or case-specific information, a manual check may be necessary. For instance, X12 requests can face difficulties with pre-authorization requirements due to standardization limitations, varying payer requirements, updates to payer policies, complex coverage rules, and challenges in coordination of benefits. These issues necessitate supplementary manual reviews or direct payer communications to ensure accuracy.

## Revenue Cycle Automation

Eligibility checks can be sent programmatically through a clearinghouse using the X12 270 format, but this method is limited because it must adhere to all payer standards using a legacy format. Additionally, the results may still lack the level of detail needed in more complex cases.

### Legacy Solutions

Directly engaging with the payer portal is one potential solution for gaining clarity; however, this method is labor-intensive and time-consuming. It requires staff to log into multiple websites, each with its own interface and navigation, complicating the process and increasing the potential for errors and delays. In response to these challenges, vendors have introduced products for revenue cycle automation, which help automate otherwise tedious tasks in the revenue cycle.

The first iteration of revenue cycle automation involved robotic process automation (RPA) solutions. RPA leveraged both frontend and API-based integrations to mirror human workflows and automate simple tasks. However, based on some public failures in the category (notably, Olive), healthcare buyers remain skeptical. Traditional RPA has faced criticism for its:

- **Brittleness:** RPA bots can break easily when interfaces change.
- **Limited adaptability:** RPA struggles with unstructured data and requires extensive manual updates.
- **Static rule-based operations:** RPA lacks the flexibility to handle complex, real-time decisions.

## AI-Enabled Solutions

Newer solutions in the revenue cycle automation category aim to overcome these pitfalls with the introduction of AI and machine learning. These technologies offer several advantages:

- **Handling unstructured data:** AI can recognize and extract relevant information from unlabeled sources, making it possible to pull both structured and unstructured data from the EHR and payer portals.
- **Dynamic adaptability:** AI/ML systems can adapt to changing interfaces and payer requirements without manual reprogramming, ensuring consistent performance despite frequent changes.
- **Complex real-time decision-making:** Advanced algorithms enable real-time decision-making, improving the accuracy and efficiency of eligibility checks.
- **Error detection and correction:** AI can detect potential errors before sending requests, reducing the likelihood of rejections and ensuring more accurate submissions.

These tools enable providers to send requests through a combination of payer portals and APIs, interacting with technical systems in the same way a human would. They can query data from APIs, navigate through web pages, extract necessary data, reason about that information, and take action on it through a web page or an API.

## Adoption Trends

Healthcare broadly has been one of the largest adopters of RPA tooling, with an estimated market size of $1.8B in 2023. According to a 2022 Becker's report, 51% of surveyed individuals reported that their organization had already adopted revenue cycle automation technologies, particularly for eligibility and prior authorization workflows.

However, despite their traction, prospective buyers should approach these technologies with realistic expectations. While AI-enabled solutions represent a significant step up in functionality over legacy RPA products, they can't yet replace humans in more complex scenarios. As a result, early adopters have found that their primary value is in reducing "button mashing" and increasing staff efficiency by enabling the redeployment of staff to other tasks.

Still, health system leaders are optimistic that the technology is steadily improving. It is anticipated that manually checking payer portals for eligibility and benefits will increasingly be supplemented or even replaced by AI-driven automation.

# Payer-Facing AI Calling

When software eligibility checks aren't enough, the alternative to engaging with the payer portal is calling the payer on the phone. In complex eligibility situations, this is the most effective way to gain clarity, but calls are operationally intensive.

## AI-Enabled Solutions

This is where payer-facing AI phone calls come into play. These AI-driven solutions can navigate IVR menus and engage with human representatives to extract detailed information and resolve ambiguities, effectively handling tasks that would otherwise require human intervention.

Thanks to rapid advancements in audio and text models, as well as the continuous improvement of LLMs for task reasoning, GenAI voice bots (such as Infinitus, Opkit, SuperDial, Amperos, and Standard Practice) are becoming increasingly adept at answering questions and dynamically generating appropriate follow-up questions.

As a result, these AI systems can manage more complex eligibility scenarios, including coordination of benefits, requests for out-of-network services, complex family coverage, and rare or high-cost treatments. If the AI fails to obtain the necessary information, some vendors also offer functionality to hand off the call to a human.

## Adoption Trends

To date, we've primarily seen adoption of these solutions by small-to-medium-sized practices. However, we anticipate increased utilization by health systems as the technology advances, given the potential for massive cost savings and efficiency gains.

## When X12 271 Isn't Enough

- **Detailed Benefits Information:** While X12 271 provides eligibility and some benefits information, it might not include detailed explanations of how benefits are calculated or specifics on copays and coinsurance rates for different services or procedures.
- **Plan Rules and Restrictions:** Specific plan rules or restrictions that apply to certain services, medications, or providers may not be fully detailed in the X12 271 response. For example, some plans require that certain procedures be performed only at specific facilities or that certain medications be sourced from specific pharmacies.
- **Alternative Treatment Options:** Information about covered alternative treatment options or less expensive service alternatives might not be provided in the X12 271 transaction.

# Prior Authorization

Prior authorization (PA) is a cost-control and utilization management process used by health insurance companies to determine whether they will cover a prescribed procedure, service, or medication, and lack of appropriate PA causes approximately 8% of all denials.

If the benefits check returns the need for a prior authorization, the provider will need to go through the requisite steps to obtain approval before providing the service or treatment. This is particularly common for expensive treatments and services like surgeries, specialty medications, and advanced imaging services. The purported goal is to verify the medical necessity and appropriateness of the care, potentially avoiding unnecessary procedures and managing healthcare costs, though some cynics might argue its real purpose is to manage utilization and deny care. In either case, it can introduce significant delays in patient care and administrative burdens for providers.

## Current Workflow

At the highest level, the PA workflow looks something like:

- The provider performs a benefits check and discovers that a prescribed treatment or medication requires approval from a patient's insurance company.
- The provider submits a request, often with necessary medical records and justification, to the payer.
- The payer reviews this information to assess its medical necessity and compliance with the patient's coverage policy.
- Following the review, the payer makes a decision to approve, deny, or request more information, which is then communicated to both the provider and the patient.

Much of the time, this happens through some combination of payer portals, faxes, and emails. In recent years, there's been a significant push from the government to standardize prior authorization and make it available via API (see CMS-0057-F), but payers have until 2027 to adopt technical standards like X12 278.

Because of the lack of standardization, health system executives we've spoken to highlight significant challenges. These include administrative inefficiencies and the slow adoption of more integrated, automated systems, resulting in prolonged decision times for approvals, denials, or requests for further information.

This has real outcomes on patient populations and on clinician burnout. A 2023 AMA Physician Survey showed that 94% of physicians report delays in care and 78% report observed instances of treatment abandonment due to prior authorization. The same survey reported that physicians and staff spend 12 hours each week, completing an average of 43 PAs per physician.

While our focus in this report is on AI in the revenue cycle, it's worth noting that there are non-AI solutions that have been adopted to attempt to manage the mayhem that is PA, namely electronic prior authorization (ePA). (See the sidebar for more on ePA vendors.)

Unfortunately, the theme across ePA solutions is that they allow payers to standardize how they process prior authorizations, while providers are forced to not only deal with differing prior auth requirements, but also contend with multiple tech stacks, portals, and steps, depending on the payer.

### Digital Prior Authorization via ePA

- Rhyme (formerly PriorAuthNow) has built its LiveAuth network, and provides a connection between registered payers and providers. It hooks into payer and utilization management (UM) vendor rules systems for prior authorization and provider EHR systems—automating steps in the prior auth process—and notifies providers and payers when manual intervention is required.

- Similarly, Availity offers a network-based approach for payers by providing a portal and integrating authorization submissions across multiple channels, then automating a portion of utilization management and prior auth tasks.
- Cohere focuses exclusively on being a workflow management platform for payers, helping automate intake and decisioning processes, and integrating with EHR systems to more easily pull clinical information. They are positioning themselves as part of a broader utilization management and steerage platform for payers.
- Surescripts offers a drug-focused ePA service that integrates into the prescription process to secure prior authorization. It runs eligibility and formulary checks to determine overall benefit status, followed by a benefit check to ascertain the cost, copay, and the necessity for prior authorization. If it is required, it coordinates the authorization process directly with the pharmacy benefit manager to resolve prior authorization requirements before the prescription reaches the pharmacy.
- CoverMyMeds, also a drug-focused ePA provider, had to take a different approach because they are not as tightly integrated into the national prescription network (Surescripts). They built their own portal and integrated with provider EHR systems, allowing them to automate much of the process, but it works after the prescription has been created, sent to the pharmacy, and rejected, and helps to resolve the rejection quickly rather than preventing it in the first place.

## AI Prior Authorization for Providers

Provider-facing software solutions for PA aim to accomplish essentially three steps:

1
- Correctly kicking off PA automatically after receiving eligibility checks indicating its necessity

↓

2
- Understanding PA requirements
- Fetching correct EHR data
- Formatting that data correctly for the PA request

↓

3
- Automating form-filling and submission

Essentially, AI-enabled PA tools get looped in as soon as eligibility checks come back. Some may be manually prompted by staff, whereas others are part of a continuous workflow directly integrated with the results of the benefits check.

Next, solutions must understand PA requirements, and these tools use a variety of methods:

- Automating phone calls directly to the payer or pharmacy benefits manager (PBM) to ascertain the specific set of conditions for approval. In some cases, this problem is so challenging that vendors like Infinitus focus entirely on gathering payer rules and status updates for other systems.
- Using ePA and clearinghouse systems to electronically receive PA requirements
- Checking PA requirements directly from EHR systems when they are already integrated with payer databases

- Scraping payer portals after logging in using provider credentials
- Predicting PA requirements based off of historical data for the procedure or medication

Based on the requirements, solutions need to determine the appropriate clinical documentation to include, then actually check that the prior authorization requirements are met. LLMs are particularly valuable in this context because they can understand relatively complex logic, then quickly search through hundreds of pages of health records to apply that logic. This allows tools to take a complex set of payer requirements into account and determine if the clinical documentation for the patient satisfies those requirements.

Assuming that the clinical documentation does match payer requirements, the next step in this automation is to provide the right set of information back to the payer, through their portals. This requires automation in form-filling with payer systems, as well as the ability to summarize unstructured data to fill out those forms in the way that a human would. Again, this is another core competency for LLMs, which can easily summarize a large corpus of data to answer specific questions about it.

Even with the best automation, though, there's no chance that every aspect of this complex workflow will be able to function without human intervention, so PA-focused tools incorporate workflow management so that operators can understand the current status for each transaction, with the option to verify or perform it manually.

Although PA is required for many payers across procedures, treatments, lab tests, imaging scans, and medications, the biggest area of development for AI-enabled prior auth workflows seems to be with prescription and specialty drugs. Vendors include Develop Health, Latent, and Tandem. The focus here is generally split between oncology drugs, specialty infusion drugs, and popular (but expensive) drugs like GLP-1s.

There are multiple reasons drug-related AI PA solutions have seen faster adoption:

- Drug PA is more standardized, whereas there are more types of procedures with more variability for acute care.
- Medications represent a high volume of PA, given prescription refills and pharmacy benefits utilization management (though acute care is growing in PA requirements).
- There is more pre-existing infrastructure with companies like CoverMyMeds and Surescripts that have established ePA on the pharma side.
- Additional incentives from drug manufacturers to make sure their drugs are easy to prescribe will often cover costs of PA.

Beyond drugs, we're seeing a number of vendors tackle a variety of health system use cases by focusing on workflow automation and integration with the EHRs and payer portals. Products like AKASA Authorization Advisor, Humata, and Thoughtful AI are incorporating both RPA to automate payer portal form submission, as well as AI to help parse through clinical documentation. Meanwhile Waystar Prior Authorization is using its direct connection with payers to check payer requirements and submit documentation, with AI powering much of the clinical summarization work.

## Adoption Trends

As interoperability across the healthcare ecosystem improves, we expect to see less value from workflow automation for navigating payer portals and filling forms and more stemming from LLMs that can query for payer rules directly, find the corresponding clinical documentation, and submit clinical justification directly back to the payer.

Interestingly, on the payer front, much of the same technology is being used to help vet prior auth submissions and streamline manual workflows. While this could create a perception of an arms race, payers' ability to correctly execute against their prior auth rules quickly and consistently is critical for patient care. On both sides, AI-enabled workflows offer the promise of faster care and less administrative overhead, but there are many hurdles to overcome before we get there.

# RCM Mid-Cycle



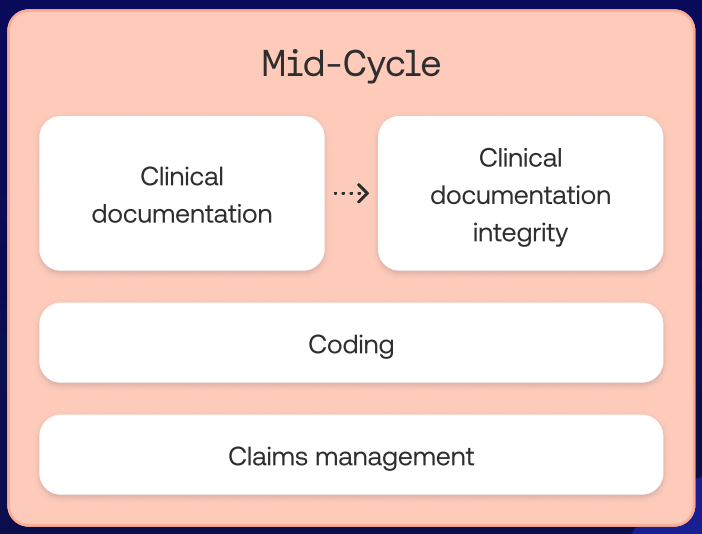


The mid-cycle of RCM essentially consists of everything that occurs from treatment to submission of the claim. It's worth noting that there is significant overlap in solutions between various steps within the mid-cycle. There are also a number of end-to-end rev cycle products, which we've included here, as claims management is one of their core competencies.

# Clinical Documentation and CDI

With treatment, comes clinical documentation. Taken during patient encounters and procedures, clinical notes are saved into the EHR for future clinical reference, and serve as the record of the encounter for legal and billing purposes.

## Current Workflow

Beyond the critical role of accurate notes in patient care, the importance of precise clinical documentation has grown steadily due to increased regulatory oversight, payer audits, and reimbursement requirements. According to AHIMA, CDI programs "facilitate the accurate representation of a patient's clinical status that translates into coded data."

As a result, there are essentially two core workflows at play:

- The initial documentation that takes place during the patient encounter
- Clinical documentation integrity (CDI), which aims to review documentation and verify it is complete and accurate

Traditionally these two steps have been executed separately and sequentially, so we've structured our discussion of solutions in the same way. In reality, though, many AI-enabled vendors in this landscape are working to merge these two steps into one.

## AI Ambient Scribes

Writing clinical documentation dates back to the early history of medical practice in ancient cultures. But the advent of EHRs and the need for increasingly precise and thorough documentation (for financial and legal purposes), has drastically increased the time and effort medical notes require.

Clinicians are increasingly spending significant amounts of "pajama time" on documenting clinical encounters, sometimes spending up to two additional hours on administrative tasks for every hour of patient care. Reports indicate that doctors often spend 1-2 hours each evening completing EHR tasks from home. This extra workload contributes to high levels of burnout, with nearly half of U.S. physicians experiencing symptoms like emotional exhaustion and depersonalization.

### Legacy Solutions

Because documentation is such a significant burden, many health systems have employed scribes to reduce clinicians' workload and improve efficiency.

- **In-room scribes:** Traditional scribes who are physically present during patient visits, taking notes and entering data into the EHR in real-time.
- **Remote scribes:** Scribes who assist with documentation remotely, often through audio or video feeds of the patient encounter.
- **Dictation/Speech-to-text apps:** Software solutions that allow clinicians to dictate their notes, which are then transcribed into text by either a human scribe or an AI-powered tool.
- **Hybrid models:** Combinations of the above methods, where a scribe might start the documentation process, and the clinician finishes it with dictation or other tools.

However, these solutions still remain imperfect because they can be costly, require ongoing training, and may not fully integrate with existing EHR systems. Additionally, even with these aids, clinicians often have to spend extra time reviewing and correcting documentation, which continues to contribute to burnout.

## AI-Enabled Solutions

With their potential to reduce costs, alleviate provider burnout, and improve patient outcomes, AI-enabled ambient scribes have become a leading category in healthcare IT—both in terms of buzz and contracts signed. For our purpose, "AI scribes" refer to systems that:

- Ambiently capture both patient consultations and clinicians' dictations.
- Convert these audio recordings into written transcripts through speech-to-text technology.
- Synthesize a clinical note from the transcript by employing artificial intelligence, often utilizing multiple large language models (LLMs) specifically tuned for this application to enhance accuracy.

Some solutions also include a "human-in-the-loop" to review the AI output before surfacing to the clinician, more closely paralleling legacy scribe solutions. Additionally, these systems typically incorporate a step to handle HIPAA compliance, which involves de-identifying protected health information (PHI) during processing and re-identifying it before the final output is delivered to ensure patient privacy and data security.

## Adoption Trends

Out of all of the technologies we've covered here, AI ambient scribes have gotten the most traction in the past few years. Health systems and large provider organizations have increasingly adopted ambient scribe solutions like Abridge, Nabla, Ambience, Nuance DAX, DeepScribe, and Augmedix after relatively short trials and proofs of concept. Where it is rolled out, clinicians have cited savings of one hour each day.

While scribes started out focused on primary care, more nuanced specialties like gastroenterology, dermatology, oncology, and hospital medicine have been able to utilize AI ambient scribes to help record clinical documentation. Additionally, we're beginning to see movement toward supporting nurse workflows with AI ambient scribes as well. Given that one of the most challenging aspects for scribes in tackling a new specialty is the medical terminology and vocabulary, this rapid progression bodes well for the underlying technology.

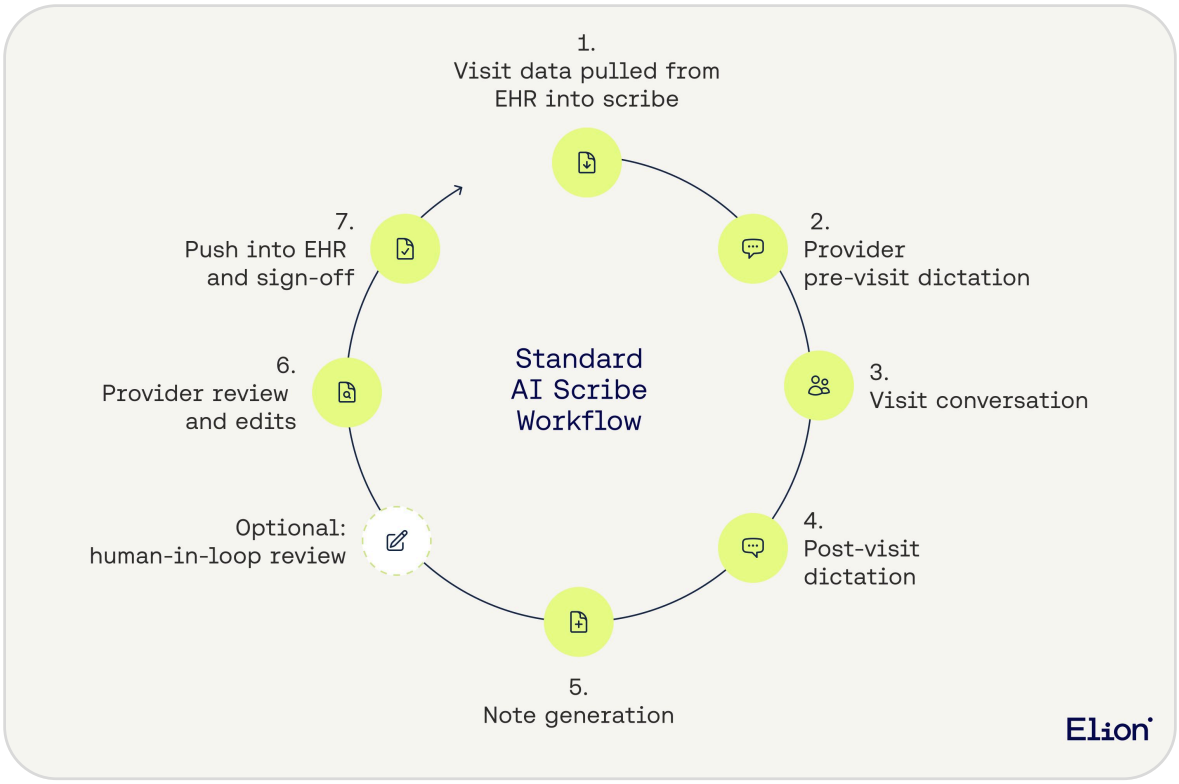


2024 Buyer's Guide to AI Scribes

# Clinical Documentation Integrity

Given the increasing imperative for complete and accurate notes, CDI is a relatively new function inserted following initial documentation and before (or in conjunction with) medical coding.

During or following a care encounter and finalization of notes by the clinician, documentation is handed off to a CDI professional. For outpatient encounters, this typically follows a visit, but inpatient care may be reviewed by a CDI specialist once every day or two during the patient's stay.

CDI specialists are trained in clinical terminology, coding guidelines, and regulatory requirements to ensure they can review documentation accurately. During the course of their work, they:

- **Select cases for review** using criteria like high-cost conditions, complex diagnoses, or cases with a high risk of audit.
- **Review clinical documentation**, looking for documentation gaps or ambiguities that could impact coding accuracy or reimbursement
- **Generate queries for clinicians** seeking clarification or additional details.
- **Receive responses from clinicians** and ensure medical records are updated for accuracy and completeness.

When performed correctly, CDI can help boost reimbursement rates and mitigate audit risks, but it requires substantial investment from both a dedicated CDI team and the clinicians seeing patients. Additionally, there is always potential for inconsistency and missed documentation improvements that aren't caught by the human eye.

## Legacy Solutions

As CDI teams developed, early entrants like Solventum (formerly 3M) 360 Encompass, Optum Enterprise CAC and CDI 3D, and Nuance CDE One realized they could also build tooling to help manage the workflows for the growing clinical documentation and coding workload. These products helped CDI teams by acting as a system of record for documentation improvement, prioritizing high-risk cases for review, tracking quality metrics, and enabling documentation improvement using complex rules engines and early NLP technology.

These products achieved broad distribution among health systems, becoming integral to the CDI process in many organizations. However, a new wave of companies built with AI from the ground up has also entered the CDI space, bringing innovative solutions leveraging cutting-edge technology from the start.

## AI-Enabled Solutions

The performance upgrade enabled by AI shifts the focus from merely prioritizing and reviewing *high-risk* clinical documentation to automatically auditing *all* clinical documentation. Generally, these tools identify documentation gaps and coding inconsistencies, making recommendations and enabling CDI specialists to improve clinical documentation more efficiently. Within this process, tools may also surface some cases to CDI specialists to validate a clinical documentation issue and submit the query to the clinician.

As previously mentioned, some of these AI-driven solutions cross over into the clinical workflow, functioning at the point of care to ensure documentation accurately reflects patient-clinician interactions, available clinical data, and best practices while supporting the assigned medical codes. For instance, Ambience, Evidently CDI, MarianaAI, and Regard integrate CDI into the clinical documentation process by proactively identifying missing clinical documentation and prompting clinicians to add it throughout the process.

Other "safety net" solutions are implemented post-encounter, reviewing clinical documentation against what has been coded and surfacing suggestions for revenue and quality opportunities. Many of these solutions include built-in functionality to suggest potential diagnoses that may have been missed based on information in the patient record. Examples include SmarterDx, Brellium, Iodine Aware CDI, 3M CDI Engage One, Optum Enterprise CAC and CDI 3D, and Semantic Health.

## Adoption Trends

Today, we're seeing legacy systems enhance their capabilities with AI integrations, while new entrants are leveraging AI from the ground up to offer innovative solutions. It's unclear which will ultimately take the lead. This dynamic landscape suggests that a blend of both legacy and new technologies may ultimately prevail, driven by the specific needs and preferences of healthcare organizations. For example:

- Newer vendors are aware of the challenges in competing with larger CDI players at health systems. Vendors like Ambience and MarianaAI are approaching from the ambient scribes/synchronous CDI angle, focusing on real-time documentation improvement during patient encounters.
- Regard, on the other hand, targets the pre-charting and clinical summarization use case for inpatient settings, helping clinicians by providing comprehensive summaries before patient visits.
- Approaching from a different angle entirely, SmarterDx runs second-pass reviews, particularly on inpatient discharge, to identify potential missed diagnoses, effectively adding a second layer of review beyond the CDI teams.
- Meanwhile, vendors like Brellium are avoiding direct competition with dominant CDI platforms by targeting smaller practices and therapy providers who may not otherwise have large CDI teams. They offer advanced CDI capabilities without the complexity and resources required for larger systems.

# Medical Coding

Typically, once the initial pass on clinical documentation has been completed, the coding process begins to capture charge codes—ICD-10 codes for diagnoses, CPT or HCPCS codes for procedures and services—along with any additional modifier codes.

## Current Workflow

The workflow for medical coding has changed dramatically over time through regulation and technological improvements, but the overall procedure remains the same. Coding involves:

1. Consolidating clinical documentation
2. Reviewing documentation
3. Assigning appropriate standardized codes
4. Entering them into the billing system to generate claims for reimbursement.

Despite extensive training, coders face increased complexity in the codes and conditions they must understand and identify. Nearly all coders now use computer-assisted coding (CAC) tools to manage their workload and accurately identify the right codes for specific procedures or conditions.

CAC technology often reads clinical notes and automatically suggests codes, shifting the coder's role to validation and ensuring the accuracy of these suggestions. Many coding teams also audit samples of each coder's output, providing feedback to correct errors and improve future performance.

While this process works well, it remains a significant source of cost (estimated at $5-$10 billion per year in the U.S.) and delay with coding turnaround times ranging from 24 hours for simple cases to a week or more for complex cases. Developments in NLP and LLMs offer the potential for faster and cheaper coding.

## AI Medical Coding

### Legacy Solutions

Early CAC systems relied heavily on rule-based NLP. These systems operated with pre-defined rules created by developers to “understand” the content of clinical documents. This approach allowed the CAC to suggest codes based on specific terms and phrases found in the documentation.

The system would search for specific keywords or phrases in the physician's notes. For example, if a note mentions a general condition like "melanoma," the system might suggest a high-level code and prompt the user to select a more specific child code. This approach required a large number of rules to attempt to handle all the scenarios in clinical documentation and had difficulty adapting to variations in language or clinical terms.

### AI-Enabled Solutions

While many AI-enabled coding tools still use some rules engines and semantic understanding, much of the work has been offloaded to more advanced NLP techniques and deep learning, which use large datasets of patient charts and their corresponding notes to implicitly teach the algorithm coding rules. These systems also improve over time by learning from user inputs, such as accepting or rejecting codes.

As a result, a new type of coding system becomes possible: While CAC systems require a human operator to verify each code, if the algorithm to predict codes becomes sufficiently accurate, a human in the loop is no longer required. Solutions that can perform coding tasks on their own without human oversight are called autonomous coding.

In addition to the specific code prediction, these algorithms can also produce a confidence score. If the autonomous coding tool finds clinical documentation that it cannot confidently process, it passes it back to traditional CAC software and human coders. In this way, most clinical documentation will be coded without human intervention, but some percentage still requires human-verified coding for now. Another practice is using autonomous coding to audit human coders, looking for potential causes of denials or downcoding.

## Adoption Trends

Three of the most dominant vendors in the autonomous coding space (CodaMetrix, Fathom, and Nym) have garnered the majority of funding and large contracts thus far. However, a growing number of upstarts are hot on their heels, including Arintra, Maverick Medical, Phare Health, RapidClaims, and Synaptec Health.

Currently, we're seeing initial adoption among low- to medium-complexity, high-volume codes: ancillary procedures like radiology and pathology, as well as primary care. More complex areas include urgent care and emergency care, where these solutions are also gaining traction.

Part of the difficulty in getting wide-scale adoption among health systems is that clinical language varies significantly across specialties. For example, the terminology used by a cardiologist can differ greatly from that of a neurologist, complicating the development and implementation of coding solutions. Contextual nuances in clinical documentation also pose challenges for even the most advanced coding systems, requiring continuous learning and adaptation for accuracy.

Autonomous coding may also struggle with inpatient facility coding due to the complexity and variability of cases, such as multiple diagnoses, procedures, comorbidities, and complications, often requiring human judgment for accurate interpretation. Additionally, multiple departments contribute to inpatient billing, adding another layer of complexity that requires coordination and oversight beyond the current capabilities of autonomous coding.

Some health system revenue cycle leaders have indicated that autonomous coding solutions don't cover enough specialties to drive sufficient efficiency to offset the cost of the system. Alternatively, some RCM teams are using multiple vendors to optimize performance across multiple specialties, since some solutions perform better in a given specialty. There is also hesitation to choose any one solution because all the major players are potential acquisition targets for dominant vendors like Optum and Solventum.

Despite these hurdles, continuous advancements in machine learning hold promise for overcoming these barriers, making it likely that autonomous coding will eventually expand its capabilities to handle more complex coding scenarios in the near future.

# Claims Management

Once the clinical documentation is complete and has been coded, providers need to create and submit claims to payers to receive payment.

## Current Workflow

A claim is essentially a detailed invoice of the coded services, along with patient and insurance information. Before sending the claim to payers, it is scrubbed to correct errors that could cause rejection or denial. Some of these edits depend on the specific service code and payer requirements, as each payer has different criteria for claims.

After in-house scrubbing, claims are typically submitted to payers through clearinghouses like Change, Availity, Waystar, TriZetto, or Experian. (Part of the reason for using these clearinghouses is that, with 1000-plus commercial payers in the US and plan-specific carve-outs, it's not always even clear which payer to route the claim to.)

These clearinghouses also perform additional scrubbing and then submit the claims electronically to payer backend systems. If claims are rejected at the clearinghouse, the provider can correct and resubmit them. Once submitted to the payer, claims are adjudicated and either paid, rejected, or denied, with the latter two requiring additional follow-up by the provider to secure payment.

## AI Claims Management

### Legacy Solutions

Historically, claims were mailed directly to payers. The introduction of HIPAA, EDI X12, and widespread EHR adoption enabled the rise of clearinghouses. As APIs and electronic infrastructure made electronic submission possible, the solutions for claims submission grew in complexity:

- EHRs and other technology vendors began integrating claims management software for providers.
- Meanwhile, clearinghouses expanded their capabilities from claims submission to include claims scrubbing, editing, and workflow management.
- Tech-oriented vendors saw opportunities with the standardization of infrastructure and built out claims management software and end-to-end revenue cycle management tools, selling into the same market as EHRs and clearinghouses.

This resulted in significant redundancies between the three solutions. Ultimately, providers still need both an EHR and clearinghouse to do all of the claims submission process, and some also leverage an end-to-RCM tool overlaid across the two. Where each step of claims submission occurs depends on a particular providers' selected vendors and preferences.

The goal of all of these claims management solutions is to help providers maximize and expedite payments, primarily by reducing the likelihood of denials. Consequently, the rules-based claims editing engines within them became increasingly complex, featuring extensive libraries of coding rules, payer policies, and regulatory guidelines that require regular updates.

However, these rules engines face limitations. The dynamic nature of coding standards, payer policies, and regulatory requirements necessitates constant updates to rule libraries. Delays in updates can result in incorrect claim edits and potential denials, as payer rules can change more quickly than scrubbers can be updated.

### AI-Enabled Solutions

Similar to the AI-enabled improvements in coding, using ML to enhance claims editing rules is a winning strategy. While it's still important to maintain a broad set of rules and a semantic understanding of claims edits, ML enables claims edit engines to implicitly learn the factors that are likely to lead to claims denial or success.

This provides a scale advantage, as clearinghouses and claims management systems that process a high volume of claims have more data to train ML-enabled rules and can detect changes to payer rules in near real-time. Conversely, upstarts with smaller claims volumes might lack the statistical significance to update rules promptly. As a result, we're seeing more partnerships between large clearinghouses and ML-oriented startups to bring this technology to market.

Major players like Solventum (in partnership with Sift) and Availity (in partnership with Anomaly) are moving beyond rules-based engines informed by multi-factored denial patterns, by teaming up with smaller, advanced AI companies. Other large clearinghouses like Change and Waystar are developing their own models. Additionally, scaling RCM companies like Candid Health and Adonis are building de novo infrastructure and data pipelines to predict denials and make claim improvements, gaining traction as they process more claims.

## Adoption Trends

While this category is still emerging, the goal is to approach all aspects of revenue cycle management with a focus on reducing denial probability. Ultimately, we expect this concept to permeate all aspects of the revenue cycle, with claims improvements starting during treatment and continuing through to payment.

## Claims Processing Flow

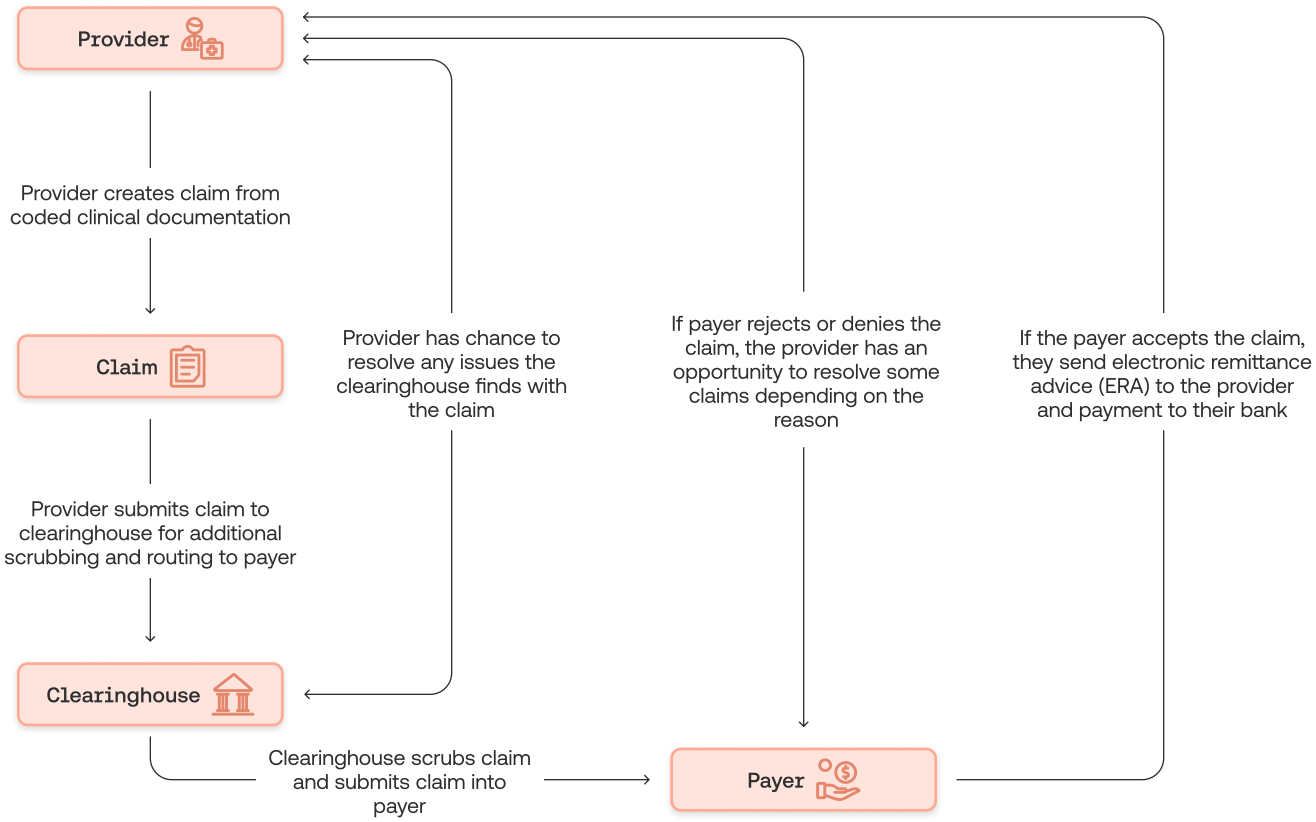

# RCM Back-End



## Back-End

Denials management

Patient billing and payment

The back-end of the revenue cycle consists of everything that happens once a claim has been submitted. In the best-case scenario, this involves only additional claims management steps such as reconciliation and payment posting, which are relatively well-handled by existing solutions. Therefore, our focus is on the more complex processes that may occur following claims submission where AI is being leveraged: denials management and patient payments.

# Denials Management

Even when everything has gone well—eligibility checks were successful, prior authorization was obtained, the clinical documentation is detailed and accurate, and the submitted claim seems right—payers can still deny the claim. In fact, this occurs for 12% of claims, with higher rates on inpatient claims. At this point, providers have the opportunity to appeal the claim, and the denials management process begins.

## Current Workflow

Larger provider organizations employ teams of revenue cycle managers, denials analysts, claims specialists, and appeals coordinators to try to mitigate claims denials. While complex in practice, the basic workflow is simple:

1. Review denial notifications from the payer and the specific reason for denial.
2. Interpret the denial reason using X12 CARC/RARC codes. These standardized codes help categorize the reasons for claim adjustments and denials. However, understanding these codes can be challenging because their meanings can vary based on factors like the payer, the type of service provided, and specific contractual agreements.
3. Review the original claim submission and supporting documentation and decide on the next steps to resolve the claim, considering the context of the denial reason, the details of the claim, and the relationship between the provider and the payer.
4. Gather additional information from clinical documentation and care teams if needed, correcting errors in the claim.
5. Submit an appeal package and detailed letter.
6. Follow up with the payer, checking appeal status until resolved.

The nuance in this process stems from managing varied policies and requirements across payers and handling complex clinical scenarios, as well as managing cost and time through the appeals process. The latter is a serious issue, as the cost to appeal a claim can vary from $25 for smaller clinics to $181 for hospitals and health systems. Additionally, while some claims with “soft denials” are possible to appeal, they may not be worth the cost, while other “hard denials” are impossible to appeal.

Beyond appealing individual denials, denials management also focuses on analyzing claims data to find upstream problems across the revenue cycle—ensuring that best practices are followed, that processes are correct, and that technology is performing as expected.

## AI Denials Management

### Legacy Solutions

Many of the same vendors for claims management also offer solutions for denials management as part of the larger RCM suite, since they already have the claims data and status available to them. These include: Change Healthcare Denial and Appeal Management, Experian Denial Management, SSI Claims Management, Waystar Denial Management, and EHR-native solutions through Epic, Cerner, and Athenahealth.

They assist in tracking work lists and preparing and submitting appeals, often including templates and automated document generation for appeal letters to ensure all necessary information is included and properly formatted.

After an appeal has been submitted, these solutions monitor communications and follow-ups with payers, tracking the status of appeals and ensuring timely actions to keep the process moving forward. Given the emphasis on preventing the source of denials, these tools also generate detailed reports on denial rates and reasons, helping to identify high-impact areas and possible root causes.

## AI-Enabled Solutions

AI tools here look to be incremental on top of existing denials management platforms. Though they can also help with managing administrative tasks like managing work lists, submitting documents, and providing time-based alerts, products like Adonis Intelligence and Rivet Claims Resolution also offer workflow automation tools to help automatically diagnose denial issues and navigate payer portals and auto-fill appeals forms.

Beyond helping with the workflow, products like Sift Denials Management predict likely denials so they can be prevented and prioritizes denied claims based on factors such as likelihood of being overturned and paid. This allows teams to focus on the most impactful cases first, optimizing time and effort. Both products leverage ML models built by unifying clinical and claims data.

For more complex appeals, GenAI is being used by tools like Crosby Health to help in drafting dynamic appeals letters and making corrections to denied claims. This includes updating coding, providing additional documentation, and correcting patient information, streamlining the appeals process.

Finally, even for claims that are paid but not in line with contract specifications, LLMs can help identify underpayments. By analyzing payment data, these tools can uncover discrepancies and ensure that providers receive the full amount they are owed. Examples here include MDClarity, Revfind, and Rivet Payer Performance.

## Adoption Trends

Most hospitals and health systems use the same vendors for both claims management and denials management, so use of AI through LLMs and workflow automation will most likely occur through adoption by the dominant players. Again, it's early in the game, but with substantial budgets and revenue on the line, we expect that health systems will invest in any technologies that can help them perform here.

# Patient Billing

Even after claims are fully paid, most patients have a portion of financial responsibility for the total cost of treatment through a copay or coinsurance. Unfortunately, almost 70% of patients' hospital bills don't get paid in full. As high deductible health plans become more popular, putting more responsibility on patients, providers are under more pressure to collect patient bills.

## Current Workflow

Even before the patient's visit, providers start by verifying coverage and benefits to determine the overall patient responsibility for treatment, and an estimate is prepared for the patient. In many cases, the patient faces substantial out-of-pocket expenses, and financial counselors are involved to help discuss payment plans, financial aid, and charity options.

After the visit, the provider submits the claim to the payer, and it is processed to generate an explanation of benefits (EOB) for the patient, as well as an electronic remittance advice (ERA) to the provider describing the remaining financial responsibility. Based on the ERA, the provider generates a billing statement for the patient, reflecting the amount owed after insurance adjustments.

The billing statement is sent to the patient via mail, email, or patient portal, along with possible payment options. The provider, or in some cases, outsourced billing services, follow up with the patient through a series of reminders and check-ins. After a longer period of non-payment, the patient is sent final notices, and then sent to collections agencies, at which point the balance is written off of the books.

## AI Patient Billing

### Legacy Solutions

One of the first innovations in helping patients make payments was web-based portals to help patients keep track of what they owed, as well as to pay their bills online, rather than having to send checks through the mail. This trend progressed as technology platforms enabled providers to start sending scheduled text message reminders and offering convenient mobile-compatible websites for viewing bills and making payments. Meanwhile, on the collections side of the house, there was very little development beyond manual workflows like phone calls and collection letters.

### AI-Enabled Solutions

We believe the category for AI patient billing is one of the most exciting areas in RCM, in part because of the consumer-facing nature of it, and the opportunity for improved engagement through personalized communication.

Like their predecessors, these payment products focus on patient engagement, bringing modern UIs and communication tools to consumers. Additionally, products like Cedar, Decoda, Flywire, Payground, Raxia, and Vantage Health are incorporating machine learning for optimizing messaging, delivery time, and message channels. Some of these products also offer billing inquiry chatbots, and startups like Ecton tackle the financial counseling part of the workflow with sophisticated chatbots and co-pilots for in-house financial counselors.

Meanwhile, other solutions like PayZen and Sift Patient Payments are more focused on optimizing payments by leveraging AI to predict (based on datasets around demographics, income levels, and credit scores) whether patients will actually pay, how much they can pay over time, and create payment plans to help patients pay off bills. In some cases, vendors are also automatically paying health systems and taking on the financial risk to collect from patients.

## Adoption Trends

Adoption of AI for patient billing is moving at a rapid pace both for smaller practices and hospitals and health systems. Traditional methods, such as web-based portals and manual collections, are quickly being augmented by AI-enabled solutions for personalized communication and engagement.

Our belief for the future is that health systems will continue to bring in technology solutions to augment or replace many of the manual processes in-house, and will eventually outsource the tail end of payments to technology tools, rather than to collections.

# The Future Role of AI in RCM

The whole RCM machine has been built around one very simple problem: *Maximizing the probability of getting paid for the clinical work you do based on everything known so far while minimizing the cost to collect that revenue.*

Historically, solving this problem has been challenging for two main reasons: there is no single source of truth encompassing all clinical data, RCM processes, claims, and claim outcomes; and, until recently, we didn't have the technological prowess to accurately predict claims outcome based off of this data.

Point solutions bring value by addressing specific issues, such as the probability of claim denials, exceptionally well. These solutions often leverage AI and other cutting-edge technologies more effectively than legacy end-to-end products, which have been slower to integrate such innovations.

However, end-to-end solutions offer their own value: they have access to data across the entire stack, enabling provider organizations to optimize the value of a claim from the outset.

The ultimate goal for RCM is to predict whether a claim will be denied or paid ahead of time, allowing for optimization at every stage:

- **Eligibility and Benefits Checking:** Can we see this patient and have it paid?
- **Prior Authorization:** What is the probability of payment for performing this procedure? Are there alternative medically appropriate clinical actions we should take?
- **Clinical Documentation:** Are we billing appropriately, and do we have all the necessary information to get paid?
- **Coding:** Do the codes align with the clinical documentation?
- **Claims Editing:** Are any changes needed to ensure the claim is correct and likely to be paid?
- **Denials Management:** Should we drop this claim or fight for it?
- **Patient Payment:** Will the patient pay? Should we offer a discount, a payment plan, or let it go to collections?

Established vendors with long-standing provider relationships and robust data pipelines have a substantial advantage trying to answer these questions. They can leverage their existing infrastructure and data to build comprehensive solutions. On the other hand, new vendors have the opportunity to incorporate cutting-edge technologies into their stack from the outset. These upstarts can build across revenue cycle steps by partnering with larger organizations that already possess access to the entire revenue cycle data.

By 2027, CMS-mandated provider-payer interoperability could significantly impact whether startups need to partner with larger vendors or can independently plug into existing data streams and grow from niche solutions. However, 2027 is still distant. Currently, larger vendors, with their established relationships and contracts, are well-positioned to endure competition while gradually incorporating AI solutions to meet customer demands.

While the question remains as to who will crack the code first—legacy, end-to-end vendors or newcomer AI point solutions—it's clear that broad adoption of AI is becoming a core part of health system strategy. The AI arms race between provider organizations and payers and regulatory forces from CMS and the FTC are strongly shaping this landscape. It's critical that health system leaders consider how they will integrate applications of AI into their organizations and build strategic frameworks for evaluating solutions across security, ROI, and integration. At Elion, we will continue to deliver research focused on the interface of AI and healthcare technology, and we aim to be your trusted source for this journey.