GeBBS®
HEALTHCARE SOLUTIONS
Forward Thinking

White Paper

Agentic AI and the Future of Autonomous Denial Recovery

Authors: Anil Reddy, Tarun Kumar Agarwal

Executive Summary

Healthcare organizations have spent years investing in denial management technology. Yet for many revenue cycle teams, it can feel like little progress has been made. You're still fighting the same battle day after day….

A denial arrives. And what should be a routine recovery suddenly becomes a race against the clock. Documentation needs to be located. Requirements need to be verified. Someone has to determine why the claim was denied in the first place. As teams bounce between payer portals, emails, and spreadsheets, the process slows under the weight of its own complexity. Then another denial arrives. And another. It never ends. You're stuck in a loop that you and your team just can't escape.

Agentic AI offers a fundamentally different approach. Rather than simply helping staff complete tasks faster, it introduces autonomous AI agents. These "digital workers" can coordinate, reason, and execute denial recovery workflows with minimal human intervention. The payoff is a more scalable recovery operation.

This white paper explores this historic shift happening in the revenue cycle now. We begin with the current denial crisis in healthcare, the evolution of denial management technology, and the emergence of Agentic AI denial recovery. The paper then pulls back the curtain on how Agentic AI denial management works, breaking down how a team of AI agents executes many of the routine tasks that have bogged down recovery teams. Finally, we'll cover the governance and oversight structures that foster trust in the AI system, as well as the technology's benefits.

As reimbursement pressures intensify and payer complexity grows, incremental improvements won't cut it anymore. A new operating model is needed. The organizations that adapt first will be best positioned to protect revenue tomorrow.

The Growing Denial Crisis

Healthcare organizations are facing a denial problem that no longer behaves like a simple administrative nuisance. It's evolved into a growing operational and financial strain that quietly spreads across the entire revenue cycle.

Claim denials continue to rise as payer policies become increasingly complex. Prior authorization requirements are expanding. Medical necessity scrutiny is intensifying. Industry estimates suggest that as many as 10–20% of claims now face initial denial, with billions of dollars in reimbursement delayed or lost each year.

What makes the problem especially frustrating is that many of these denials are technically recoverable. The challenge is, many organizations don't have the operational bandwidth to pursue recovery efficiently. That's where denial management often begins to break down.

After a denial occurs, teams often find themselves stuck in a chain reaction. An analyst begins digging into the root cause while someone else searches for missing documentation. Meanwhile, another department may need to weigh in before an appeal can even begin. As work moves from queue to queue, the account ages. And all the while, the reimbursement clock keeps ticking.

And despite years of investment in automation, many organizations still feel like they're fighting these battles one denial at a time. How is that possible when automation was supposed to improve efficiency?

The answer lies in how denial management technology has evolved and where its limitations began to surface.

How We Got Here: The Evolution of Denial Management

The denial management systems of today didn't appear overnight. They evolved gradually in response to growing payer complexity and mounting administrative pressure. But each phase also introduced new limitations that eventually pushed the industry toward the next stage of innovation.

Phase 1: Manual Denial Handling

In the early days of denial management, nearly everything was handled manually. Analysts reviewed Explanation of Benefits (EOBs), interpreted CARC and RARC codes, searched patient records, drafted appeal letters, and tracked outcomes through spreadsheets.
The process was slow and heavily dependent on human expertise. Experienced staff often became the "tribal knowledge" holders of payer behavior and denial resolution strategies. That worked for a while. But as denial complexity increased, scalability became a growing problem.

Phase 2: Rule-Based Automation

To improve efficiency, healthcare organizations introduced rule engines and robotic process automation (RPA). These systems helped standardize workflows and automate repetitive tasks.
While the gains were meaningful at first, over time, the limitations became difficult to ignore. Rule-based systems are rigid by design. Every denial scenario requires predefined logic, and healthcare rarely stays still long enough for static workflows to keep pace. Payer requirements change constantly. New edge cases emerge. Organizations often find themselves trapped in a cycle of continuously updating rules just to maintain operational stability.

At its core, traditional automation lacks contextual reasoning. It can follow instructions, but it can't truly understand meaning.

Phase 3: Generative AI Assistants

Generative AI introduced another major leap forward. Suddenly, systems could summarize records, extract documentation, and even help draft appeal letters.
That said, most generative AI solutions still function primarily as assistants or copilots. They help humans complete tasks faster, but they generally rely on us to initiate workflows and make operational decisions.

And that distinction matters. Because at the root of the denial management problem is a lack of autonomous orchestration. Too much manual intervention is involved. Content generation doesn't solve that problem.

Phase 4: Agentic AI–Driven Autonomous Denial Recovery

Agentic AI represents the next evolution.
Instead of relying on isolated automation tools or passive AI assistants, Agentic AI introduces autonomous systems. These systems can reason, collaborate, plan, and execute multi-step denial recovery workflows with minimal human intervention. What's more, Agentic AI continuously learns from outcomes. It gets smarter.

The shift is subtle but powerful. Rather than simply assisting teams, the system now actively manages denial recovery operations itself.

What Is Agentic AI in Post-Denial Management?

One of the simplest ways to think about Agentic AI is as a team of specialized autonomous digital workers (or robots), each responsible for different parts of the denial recovery process. That idea becomes easier to grasp when compared to how people work in real life.
In healthcare operations, no single person handles every responsibility. One employee may focus on payer communication. Another may gather clinical documentation. Another oversees appeals strategy. Each person contributes specialized expertise to

move the process toward resolution. Just like humans create better outcomes when working as a team, so do agents. And Agentic AI systems mirror that team structure.

Instead of relying on one massive AI system to handle every denial scenario, organizations deploy multiple specialized agents that work together. The employee gathering clinical documentation? That becomes an agent. The one responsible for appeals strategy? An agent handles that now.

And just like in real life, someone still must oversee the entire operation. In this case, that role belongs to the orchestration layer—an agent that coordinates workflows, assigns tasks, escalates issues, and helps move denials toward resolution.

In many ways, the system behaves less like a traditional software tool and more like a synchronized operational team. And synchronization is where the model becomes especially powerful.

In traditional workflows, delays naturally occur between handoffs because human teams operate asynchronously. When one analyst finishes a task, the next person may not review it until hours or days later because of competing priorities. Autonomous agents don't have that limitation.

As long as the necessary data is available, multiple agents work simultaneously. Clinical evidence retrieval no longer has to wait for manual handoffs. Agents can start retrieving potentially relevant clinical evidence while payer requirements are being analyzed, refining the search as the requirements become clearer. The same goes for appeal preparation. As one agent continues to gather supporting documentation, another can prepare the appeal, allowing multiple steps to progress at once.

The result is a dramatically faster denial recovery operation capable of adapting, learning, and improving over time—with minimal human involvement.

How Agentic AI Denial Management Works

At first glance, Agentic AI denial management can sound almost futuristic. Multiple AI agents collaborating autonomously to recover denied claims? It's easy to assume the process must be incredibly complicated behind the scenes.

In reality, the core concept is surprisingly intuitive. A team of specialized agents works together to achieve a desired outcome. Here's what that looks like in denial management.

Denial Intake Agent

Every denial recovery workflow begins with understanding why the claim was denied.

That responsibility falls to the Denial Intake Agent. Its job is to ingest denied claims from ERA files, payer portals, clearinghouse feeds, and billing systems, then classify and interpret the denial details associated with each claim.

At first, that may sound fairly straightforward. But denial interpretation is often far messier than organizations would like to admit.

Different payers frequently describe similar denial issues in completely different ways. One payer may reference "medical necessity." While another may cite "lack of conservative therapy" or use an entirely different denial code altogether. Human analysts often recognize these patterns instinctively because they understand the broader context. However, traditional automation systems tend to struggle because they operate within predefined rules.

The Denial Intake Agent bridges that gap by using natural language processing (NLP) to interpret and standardize denial information. With this, it recognizes when different denial codes, descriptions, or payer terminology represent the same underlying denial themes across disparate payer language.

Root Cause Analysis Agent

Once the denial is identified, the next challenge is determining what caused it. And that's not always as obvious as the denial code suggests.
A denial categorized as "medical necessity" could stem from missing documentation, coding inconsistencies, authorization gaps, modifier issues, or payer-specific policy requirements. That's where the Root Cause Analysis Agent comes in. It's designed to move beyond surface-level denial descriptions and identify the underlying issue.

To do this, the agent analyzes historical denial patterns, documentation gaps, eligibility mismatches, coding behavior, and payer policy violations using large language model (LLM)-driven analysis.

This represents one of the major differences between traditional automation and Agentic AI. Traditional workflows often react to denial labels. Agentic systems attempt to understand the broader operational context surrounding the denial itself.

Clinical Evidence Retrieval Agent

Once the root cause is identified, the system must gather the supporting documentation required for appeals. That responsibility belongs to the Clinical Evidence Retrieval Agent.
This agent autonomously retrieves physician notes, operative reports, lab results, imaging findings, medication histories, and prior treatment documentation from clinical systems and EHR environments.

Importantly, the system is not simply keyword matching.

Modern agentic systems can understand unstructured clinical narratives. Instead of searching only for exact terminology, an agent may recognize medically relevant evidence buried within physician documentation. For example, it may identify evidence of neurological deficits or failed physical therapy—even when those concepts aren't explicitly stated. This capability helps identify clinically relevant evidence more consistently, while also reducing manual chart review workloads.

Payer Policy Intelligence Agent

One of the biggest frustrations in denial management is that payer requirements rarely stay static. Policies evolve. Medical necessity requirements change. And often, those changes happen faster than organizations can adapt.

The Payer Policy Intelligence Agent is designed for just this purpose. It continuously monitors and interprets payer behavior, policy updates, LCD and NCD changes, and denial trends.

Over time, the system builds institutional payer intelligence by learning from operational feedback and denial outcomes. That capability becomes especially

important in large healthcare organizations where payer variation can quickly overwhelm teams trying to keep up with changing requirements.

Appeal Strategy Agent

Not every denial should be handled the same way. Some claims should be appealed aggressively. Others need to be corrected and resubmitted. Some may not be financially viable to pursue at all.

The Appeal Strategy Agent helps determine the optimal recovery pathway for each denial scenario.

To make those decisions, the agent evaluates multiple factors simultaneously, including:

- Denial type
- Payer behavior
- Historical appeal success rates
- Claim value
- Documentation quality
- Filing deadlines

In many ways, this agent functions like an operational decision engine designed to maximize recovery efficiency and reimbursement, rather than blindly automate tasks.

Autonomous Appeal Generation Agent

Once there's an optimal strategy, the next step is executing the appeal.
The Autonomous Appeal Generation Agent creates complete appeal packages that may include medical necessity justifications, policy references, evidence summaries, and physician narratives.

And critically, the agent can tailor its language based on the payer.

An appeal prepared for Medicare may require very different wording than one prepared for a commercial payer or managed care organization. The agent adapts accordingly, helping create more contextualized and payer-specific appeals.

Learning and Optimization Agent

The last specialist digital employee involved is the Learning and Optimization Agent. Its job is to continuously improve the system and refine its decision making over time.

It learns from:

- Appeal outcomes
- Payer responses
- Successful recovery patterns
- Denial trends
- Operational bottlenecks

That continuous feedback loop is critical because denial management isn't a static process. Payer behavior evolves constantly. Agentic systems are designed to evolve alongside it.

Workflow Orchestration Agent

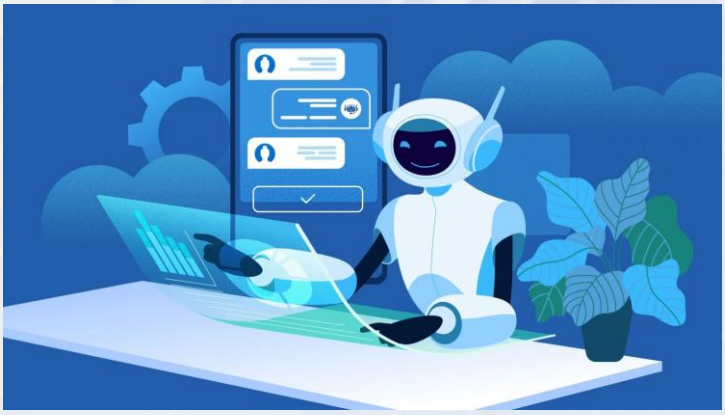

Overseeing all these agents is the Workflow Orchestration Agent—the conductor coordinating the entire denial recovery ecosystem.

As alluded to in the last section, its role is to manage the process. And like any good manager, it will assign tasks, prioritize denials, monitor SLAs, coordinate workflows between staff (or in this case, agents), and ensure the recovery process moves efficiently along.

Without orchestration, the system would simply become a collection of disconnected AI tools. The orchestration layer is what transforms those individual capabilities into a synchronized operational engine.

Governance, Compliance, and Human Oversight

Whenever AI enters healthcare operations, one question inevitably follows…can the system be trusted?

It's a fair concern. In denial management, the stakes are high. Claims involve protected health information, payer compliance requirements, reimbursement risk, and operational decisions that directly impact financial performance. The idea of autonomous AI agents making decisions inside those workflows can understandably create some hesitation.

But Agentic AI denial management is not designed to operate without guardrails. In reality, governance, explainability, compliance controls, and human oversight are foundational components of the architecture. One of the biggest misconceptions of autonomous AI systems is that they function like black boxes—as if they make invisible decisions behind the scenes. In well-designed healthcare AI environments, that simply cannot happen.

Every action performed by an agent can be logged, audited, and traced back to the reasoning and evidence used to support the decision. If an appeal is generated, the system can identify:

- The denial that triggered the workflow
- Retrieved documentation
- Referenced payer policies
- Why a particular recovery strategy was recommended

That level of explainability becomes especially important in healthcare environments where audit readiness and compliance oversight are critical.

Healthcare organizations adopting Agentic AI systems must still maintain strict adherence to HIPAA requirements, data security standards, and payer regulations. That means access management, encryption, workflow permissions, secure API integrations, PHI protection, and human approval checkpoints remain essential parts of the environment.

Where Do Humans Fit in the Process?

Despite the growing sophistication of autonomous agents, organizations still need experienced professionals to review highly complex cases. In many ways, the role of denial teams evolves rather than disappears.

Instead of spending hours manually chasing documentation or navigating repetitive workflows, staff can focus their attention where judgment matters most:

- Complex appeals
- Escalated cases
- Payer disputes
- Policy interpretation
- Strategic recovery decisions

People often assume AI in healthcare is primarily about replacing humans. But in practice, many organizations find the opposite. Denial analysts increasingly move from being task executors to escalation specialists and strategic decision-makers.

Business Benefits

The operational benefits of Agentic AI denial management extend far beyond simple task automation. Here are some of the most intriguing.

Reduce Denial Recovery Timelines from Weeks to Hours

Agentic AI eliminates the manual handoffs and sequential task execution inherent in traditional denial recovery. It replaces that with synchronized workflows, where multiple agents can work simultaneously—enabling documentation retrieval, appeal

preparation, and similar tasks to occur in parallel. The result is sheer speed. In some cases, denial resolution timelines shrink from weeks to hours.

Scale Without Linear Staffing Growth

As denial volumes rise, many organizations struggle to scale through staffing alone. Agentic AI introduces a different model. Instead of relying entirely on headcount growth, organizations can leverage AI agents to autonomously handle much of the workflow. This allows denial teams to manage larger claim volumes without the same level of operational strain or administrative backlog.

Lower Write-Offs and Reduce Revenue Leakage

Many denials are recoverable. The challenge is knowing which claims are worth pursuing, and acting before recovery opportunities disappear. By analyzing denial reasons, payer behavior, documentation quality, claim value, and historical outcomes, Agentic AI helps identify recoverable claims more accurately. This allows organizations to prioritize the right accounts and capture revenue that may otherwise fall through the cracks.

Strengthen Data-Driven Decision-Making

Agentic AI provides greater visibility into the denial recovery process. It does this by continuously analyzing data points, such as payer responses and recovery outcomes. Over time, organizations gain clearer insight into where denials originate, which payers create the most friction, and which recovery strategies produce the strongest financial outcomes. This data ultimately helps leaders make more informed operational and financial decisions.

Prevent More Denials Over Time

Agentic AI may eventually help healthcare organizations move beyond denial recovery and toward denial prevention. How? The system learns from recurring

documentation gaps, payer behavior, and workflow inefficiencies. With this knowledge, it could self-adjust to prevent many denials in the first place.

Redefining the Future of Denial Management

Being overwhelmed by denials no longer needs to be the status quo. With Agentic AI, a new future can be forged.

When a denial enters the system, work can begin immediately. No time needs to be wasted on handoffs and tracking down information. AI agents got you covered. One agent gathers supporting evidence while another reviews payer requirements.

By the time a human recovery specialist opens the case, much of the investigation has already been completed. Fewer claims linger unresolved. Fewer recovery opportunities are missed. Instead of spending the day chasing information, teams spend more time applying expertise where it matters most.

That is the promise of Agentic AI in denial management. Not simply doing the same work faster, but creating a revenue cycle operation that learns, adapts, and keeps moving forward, even as payer complexity continues to rise.

The future described in this paper is closer than you may realize. At GeBBS, we've already developed the Agentic AI denial management solution presented in this white paper. It's designed to autonomously analyze denials, retrieve supporting documentation, optimize recovery strategies, and accelerate appeal workflows. Rather than adding another tool for staff to manage, why not change the operating model itself? The benefits may just be transformative. Not only will you reduce denials, but you'll also recover revenue faster and improve efficiency. Want to know more? Contact us today to talk with a specialist.