GeBBS®
HEALTHCARE SOLUTIONS
Forward Thinking

Agentic AI Platform for Autonomous End-to-End Revenue Cycle Management - From Automation to Autonomy

White Paper

Summary

Healthcare revenue cycle management (RCM) remains one of the most complex and labor-intensive operational domains in modern healthcare. Fragmented systems, manual workflows, payer variability, and compliance pressures contribute to inefficiencies, claim denials, and revenue leakage.

This white paper introduces an Agentic AI Platform for Autonomous End-to-End Revenue Cycle Management—a computer-implemented, intelligent, and self-optimizing RCM system powered by coordinated AI agents, governed by a Model Context Protocol (MCP), and supported by a shared cross-agent knowledge fabric.

The platform autonomously executes the entire RCM lifecycle—from patient onboarding to final payment reconciliation—while continuously learning and improving performance through outcome-driven adaptation. By replacing static workflows with collaborative, reasoning-driven AI agents, the system significantly improves coding accuracy, reduces denials, accelerates cash flow, and minimizes manual workload.

1. Introduction

Healthcare providers face mounting pressure to improve financial performance while maintaining regulatory compliance and delivering high-quality patient care. Traditional RCM systems rely heavily on:

Manual data entry and interpretation

Rule-based automation with limited flexibility

Siloed departmental workflows

Retrospective error detection

These limitations lead to delayed reimbursements, preventable denials, and operational inefficiencies.

The proposed Agentic AI RCM Platform addresses these challenges by introducing a multi-agent, autonomous, and adaptive architecture that performs real-time reasoning across the revenue cycle.

2. Limitations and Challenges in Current RCM Workflow

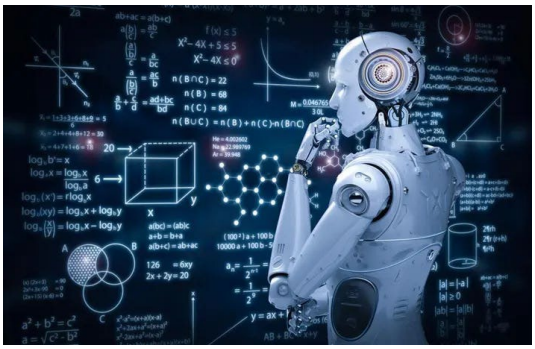

Despite technological investments, current RCM solutions suffer from fundamental structural limitations:

2.1. Fragmented Workflow & Data Silos

- Disconnected systems across scheduling, coding, billing, and collections
- Lack of shared context across departments
- Redundant data entry and reconciliation

2.2. High Denial Rates and Revenue Leakage

- Missing authorizations
- Incorrect coding
- Documentation gaps
- Payer-specific policy complexity

Studies show that claim denial rates can exceed 10–15% across providers, significantly impacting revenue cycles.

2.3. Manual Effort and Operational Cost

- 40–60% of RCM activities remain manual
- Significant labor cost for repetitive tasks
- Human error and inconsistent decision-making

2.4. Static Rule-Based Automation

- Traditional RPA cannot adapt dynamically
- Hard-coded workflows break with payer rule changes
- Limited learning capability

2.5. Poor Clinical-Financial Alignment

- Clinical documentation often not optimized for billing
- Coding teams interpret charts differently
- Lack of CDI (Clinical Documentation Integrity) automation

2.6. Limited Intelligence and Predictive Capability

- Existing analytics are retrospective
- No real-time decision optimization
- No autonomous corrective feedback loops

3. How Agentic AI Resolves Existing RCM Challenges

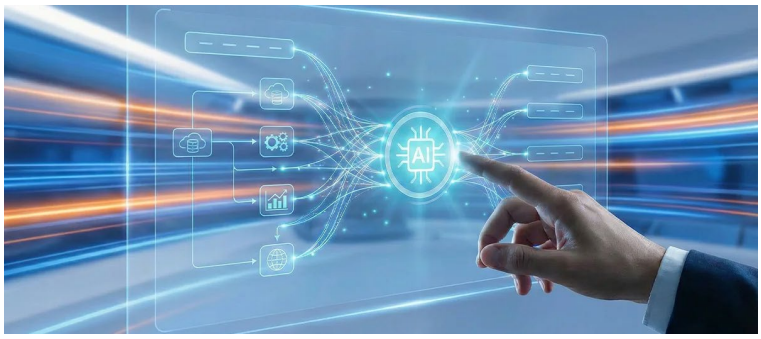

Traditional RCM systems struggle with fragmented workflows, manual handoffs, static rule engines, and lack of real-time intelligence. Agentic AI addresses these limitations through a multi-agent, autonomous, and self-learning architecture where specialized agents collaborate, share context, and continuously improve outcomes. This approach replaces siloed operations with coordinated intelligence across the entire revenue cycle—from patient access to payment reconciliation—enabling proactive issue prevention, faster decision-making, and end-to-end optimization.

3.1. What is Agentic AI in RCM?

Agentic AI in RCM refers to a network of collaborative intelligent agents, each specialized in a domain function such as eligibility, coding, billing, denial management, or payment posting. These agents operate independently yet cohesively under a shared orchestration and governance layer.

Key characteristics include:

- **Autonomous Decision Making**
  Each agent independently interprets context, applies domain logic, and executes tasks such as coding a chart, generating an authorization request, or initiating an appeal—without manual intervention.

- **Context Sharing Across Agents**
  Agents operate on a shared memory fabric where clinical data, payer policies, coding rationales, and denial outcomes are accessible system wide. This ensures that decisions made in one stage (e.g., authorization) inform downstream actions (e.g., coding and billing).

- **Continuous Learning from Outcomes**
  The system continuously captures results such as claim acceptance rates, denial causes, and reimbursement variances. These outcomes are fed back into agents to refine future decisions and reduce recurring errors.

- **Coordinated Task Orchestration**
  Agents communicate through an orchestration layer using intent-based events (e.g., “Denial_Detected,” “Authorization_Required,” “Recoding_Needed”). This enables seamless end-to-end workflow automation across the RCM lifecycle.

- **Human Escalation When Required**
  For complex, ambiguous, or high-risk scenarios, agents escalate tasks to human experts based on confidence thresholds, compliance policies, or predefined governance rules.

3.2. Core Capabilities of Agentic RCM

A. Contextual Intelligence

- Agentic AI leverages large language models (LLMs) combined with retrieval-augmented generation (RAG) to interpret both structured and unstructured healthcare data.
- It reads physician notes, operative reports, lab results, and radiology findings to understand clinical context.
- It retrieves payer-specific rules, authorization policies, LCD/NCD requirements, and coding guidelines in real time.
- It maintains persistent patient-level and claim-level context across the entire lifecycle, ensuring continuity from scheduling to final payment.
- This eliminates information silos and ensures every decision is made with complete, up-to-date context.

B. Autonomous Workflow Execution

In Agentic RCM, workflows are not manually triggered—they are automatically orchestrated by agents. When a chart is completed, the coding agent is triggered automatically. Once coding is completed, the billing and claim creation agents take over. The claim submission agent validates and submits claims to clearinghouses without human intervention. This enables straight-through processing of claims with minimal manual touchpoints and significantly faster throughput.

C. Real-Time Decisioning

Agentic AI enables real-time validation and decision-making at every stage of the revenue cycle:

- Dynamic payer policy validation ensures compliance with evolving payer rules before claim submission.
- Authorization checks before services prevent non-covered procedures from being scheduled or performed.

- Coding accuracy checks before submission ensure alignment with documentation and payer edits.
- By shifting validation upstream, the system prevents downstream denials and rework.

D. Self-Learning Feedback Loop

A key differentiator of Agentic AI is its closed-loop learning system. Denial outcomes are automatically analyzed and categorized.

Insights are fed back into: coding logic (to avoid incorrect code selection), CDI workflows (to improve documentation quality), payer rule interpretation (to refine authorization and billing compliance). Over time, the system learns from every claim, reducing recurring errors and continuously improving performance.

E. Human-in-the-Loop Governance

- While Agentic AI automates most workflows, it operates within a governed human oversight framework:
- Complex or low-confidence cases are flagged for expert review.
- Escalation thresholds are based on AI confidence scores, compliance risk, or financial impact.
- All agent decisions include explainability artifacts (rationales, source references, confidence scores), enabling human auditors to validate outcomes.

This ensures that autonomy is balanced with compliance, transparency, and clinical safety.

4. Architectural Foundations

4.1. Agentic AI Orchestration Framework

The platform is built on a distributed agent architecture where each agent:

- Performs a specialized RCM function
- Reasons overstructured and unstructured healthcare data
- Collaborates with other agents using shared context
- Learns from outcomes to improve future decisions

This replaces rigid workflow pipelines with dynamic, intent-driven orchestration.

4.2. Model Context Protocol (MCP)

The MCP layer is the backbone of the system and provides:

- Standardized context exchange
- Secure tool and data access
- Governance and approval controls
- Intent-based agent communication
- Interoperability across systems and agents

All inter-agent communication flows through MCP, ensuring traceability, compliance, and explainability.

4.3. Cross-Agent Knowledge and Memory Fabric

A shared intelligence layer that stores:

- Clinical documentation
- Coding decisions and rationales
- Payer policies and authorization rules
- Denial histories and appeal outcomes
- Payment patterns and contract terms

Agents read from and write to this shared memory, enabling consistent reasoning across the entire revenue cycle.

5. Autonomous RCM Agents

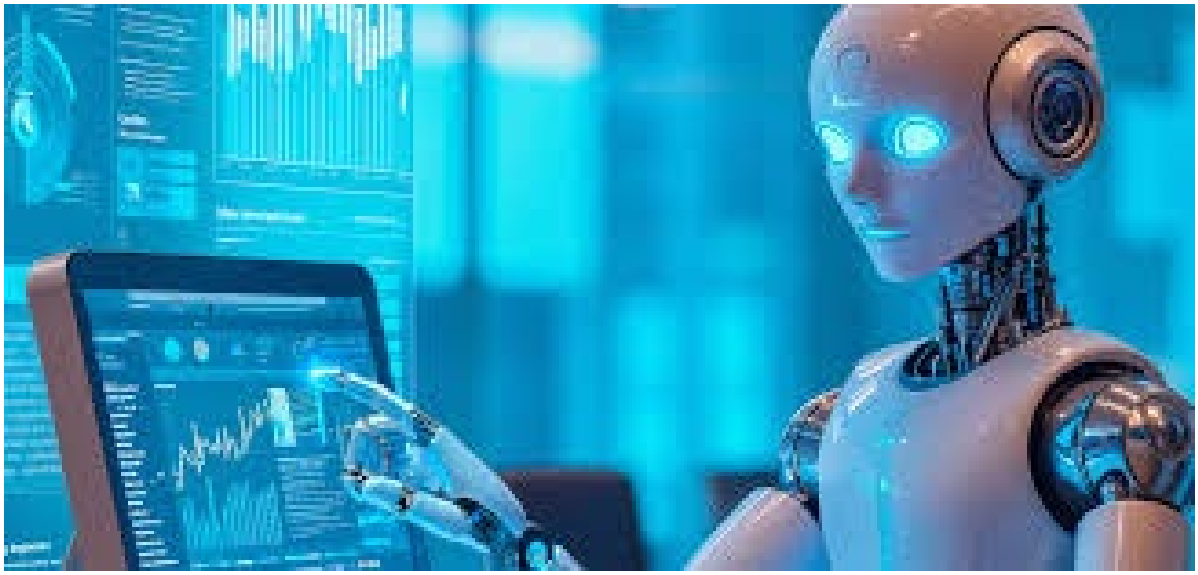

5.1. Patient Registration & Scheduling Agent

This agent autonomously manages patient onboarding by capturing demographics, insurance details, and patient intent through conversational interfaces such as voice, chat, or digital forms. It validates and corrects incomplete or inconsistent data in real time while optimizing appointment scheduling based on provider availability, specialty needs, and patient preferences. The agent outputs structured patient records, scheduling metadata, care pathway classification, and confidence scores for data completeness.

5.2. Insurance Verification & Authorization Agent

This agent ensures financial readiness before care delivery by verifying eligibility, benefits, and coverage details while interpreting complex payer policies. It determines prior authorization requirements, generates compliant authorization requests, and estimates patient financial responsibility. The agent proactively identifies authorization risks and initiates corrective actions to prevent downstream denials and delays.

5.3. Medical Coding & Auditing Agent

This agent performs real-time, AI-driven coding by analyzing clinical documentation using advanced reasoning models to assign accurate CPT, ICD-10, PCS, and HCC codes. It continuously audits coding decisions for compliance, detects documentation gaps, and produces explainable coding rationales. The agent can be

automatically re-invoked by the denial workflow to perform targeted recoding and improve claim outcomes.

5.4. Clinical Documentation Improvement (CDI) & Charge Capture Agent

This agent ensures documentation completeness and revenue integrity by identifying missing or under-documented clinical elements during care delivery. It generates compliant physician queries, recommends documentation enhancements, and aligns charge capture with clinical evidence. By operating concurrently with clinical workflows, it prevents revenue leakage before claim submission.

5.5. Denial Management & Appeals Agent

This agent autonomously manages denied claims by interpreting EOBs and payer denial narratives to identify root causes such as coding, authorization, or documentation issues. It generates tailored appeal strategies with supporting evidence, triggers recoding or CDI workflows when necessary, and continuously learns from historical appeal outcomes. This enables proactive denial prevention and improved recovery rates over time.

5.6. Payment Reconciliation Agent

This agent completes the revenue cycle by interpreting ERA and remittance data to reconcile expected versus actual payments. It detects underpayments, contractual discrepancies, and posting errors, and initiates recovery workflows automatically. The agent provides real-time financial visibility, ensuring accurate revenue recognition and closure of the revenue loop.

6. Cross-Agent Coordination and Communication

Unlike traditional linear workflows, the platform uses intent-based, event-driven collaboration.

Key Communication Events

- Denial_Detected
- Authorization_Risk_Identified
- Coding_Clarification_Requested
- Documentation_Insufficiency_Detected
- Underpayment_Pattern_Observed

Agents respond to these events collaboratively, using shared context and reasoning to resolve issues in parallel and in real time.

7. Overall Benefits of Agentic AI in Revenue Cycle Management (RCM)

Agentic AI introduces a fundamentally new operating model for healthcare revenue cycle operations—one that is autonomous, collaborative, and continuously improving. Below is a detailed expansion of the key benefit areas and how they translate into measurable business value.

7.1. Financial Impact

- **Reduced Denial Rates**

  Agentic AI proactively prevents denials by validating eligibility, authorization, coding, and payer policy requirements before claim submission. Denial agents continuously analyze payer responses and feed insights upstream to coding, CDI, and authorization agents—eliminating recurring error patterns. This shifts denial management from reactive correction to preventive intelligence.

- **Faster Reimbursement Cycles**

  By automating claim preparation, scrubbing, submission, and follow-ups, Agentic AI significantly reduces billing cycle time. Claims are submitted faster, errors are corrected earlier, and AR follow-ups are prioritized using AI-driven aging and value scoring. This accelerates cash flow and reduces Days in AR.

- **Increased Clean Claim Rate**

  Through real-time validation and intelligent claim assembly, the system improves first-pass acceptance rates. Coding, documentation, and payer compliance are validated before submission, resulting in a higher clean claim rate and reduced rework.

7.2. Operational Efficiency

- **40–60% Reduction in Manual Workload**

  Repetitive tasks such as eligibility checks, coding, claim submission, payment posting, and denial classification are automated through specialized agents. Staff focus shifts to exception handling and high-value cases.

- **Faster Processing Time**
  Agentic systems operate in near real-time and in parallel across workflows. For example, eligibility verification, coding, and authorization checks can run simultaneously, reducing cycle times from days to minutes or hours.

- **Scalable Operations without Linear Headcount Growth**
  Agentic AI scales horizontally—processing higher claim volumes without proportional increases in staff. This allows organizations to grow revenue without increasing operational cost base.

7.3. Accuracy and Compliance

- **Reduced Coding Errors**
  Medical Coding Agents use large language models and clinical reasoning to align documentation with coding guidelines, payer edits, and risk adjustment rules. Continuous auditing ensures errors are detected and corrected in real time.

- **Automated Policy Validation**
  Authorization rules, LCD/NCD policies, NCCI edits, and payer-specific requirements are dynamically retrieved and applied during claim preparation—ensuring compliance before submission.

- **Full Audit Trails and Explainability**
  Every decision made by an agent includes:
  - Source data references
  - Reasoning traces
  - Confidence scores
  - Policy citations

  This provides end-to-end auditability, supporting compliance reviews and payer audits.

7.4. End-to-End Visibility

- **Real-Time Dashboards**
  Agentic RCM platforms provide live operational dashboards showing:
  - Claim status across lifecycle
  - Denial trends
  - AR aging
  - Authorization backlog

- **Predictive Revenue Insights**
  AI models forecast:
  - Expected reimbursement

- Denial likelihood
- Cash flow projections
- Underpayment risk

This enables proactive financial decision-making.

- **Root Cause Analytics**
  Denials, underpayments, and delays are traced back to upstream causes (documentation gaps, coding errors, authorization issues). This enables continuous process improvement across departments.

7.5. Improved Patient Experience

- **Transparent Billing**
  Patients receive accurate cost estimates upfront and clear billing statements with explanations of charges, insurance coverage, and out-of-pocket responsibility.

- **Faster Claims Resolution**
  Automation reduces claim turnaround time, resulting in faster adjudication and fewer billing surprises for patients.

- **Personalized Financial Engagement**
  AI agents can offer:
  - Tailored payment plans
  - Proactive billing reminders
  - Conversational support via chat/voice agents

  This improves patient satisfaction and increases collections.

7.6. Continuous Learning System

- **Denials Continuously Improve Upstream Processes**
  Each denial is analyzed and converted into a learning signal. These signals update:
  - Coding strategies
  - Documentation requirements
  - Authorization workflows

  Over time, the system eliminates recurring denial patterns.

- **AI Agents Evolve Over Time**
  Agents refine their decision thresholds, reasoning logic, and workflow routing based on outcomes such as approval rates, reimbursement variance, and appeal success.

- **Knowledge Graph Expands Dynamically**
  The system continuously builds a knowledge base of:

- Payer rules
- Coding rationales
- Denial patterns
- Contract terms

This evolving intelligence improves performance across all future claims.

8. Benefits of Agentic AI based RCM to Healthcare Organizations

For healthcare organizations (health systems, hospitals, physician groups, RCM enterprises, and MSOs), Agentic AI–driven RCM is not just an automation tool—it is a strategic digital revenue infrastructure that delivers:

- Higher revenue
- Lower cost
- Better compliance
- Scalable operations
- Data-driven leadership insights
- Superior patient financial experience

It transforms RCM from a back-office function into a strategic growth engine.

9. How GeBBS is using Agentic AI in RCM?

GeBBS applies an agentic AI framework across the RCM lifecycle, linking patient registration, insurance verification, medical coding, CDI, charge capture, claim submission, denial management, and payment reconciliation. LLM and ML agents extract ICD/CPT codes, validate eligibility and authorizations, predict risks, and enable direct-to-bill workflows. Claims are scrubbed and submitted automatically, while denial agents analyze EOBs, identify root causes, trigger recoding or CDI, and generate appeals. Continuous feedback loops improve coding accuracy, reduce denials, accelerate reimbursement, and provide real-time financial visibility.

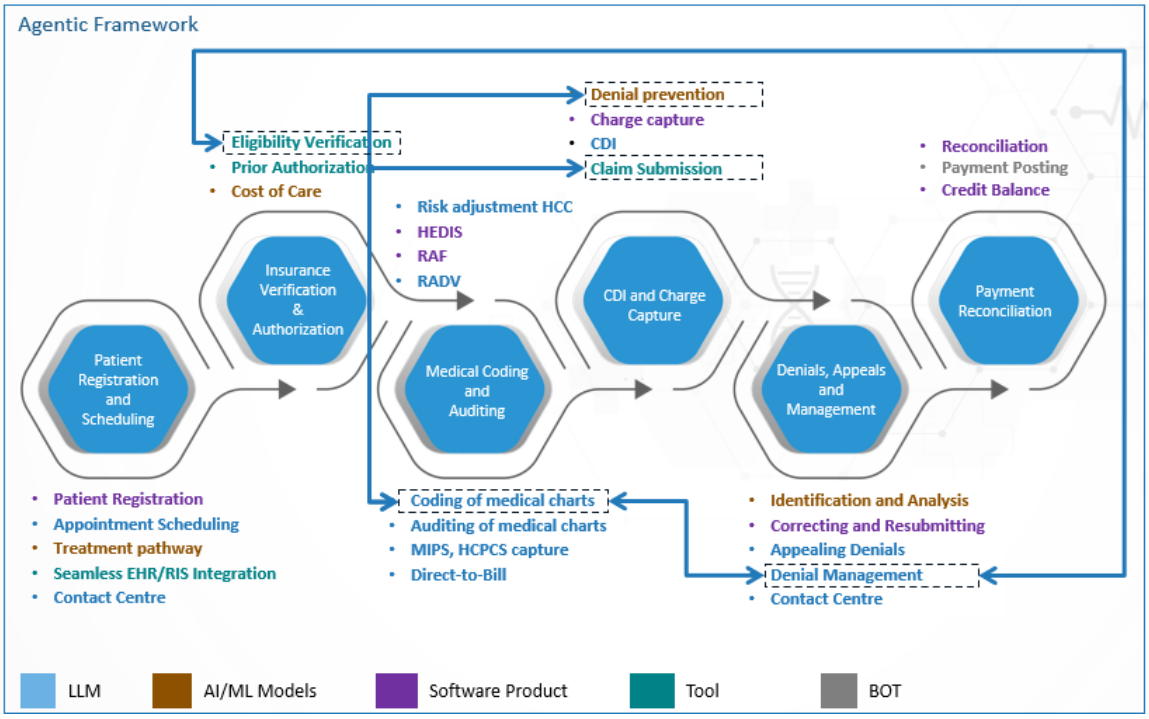

10. Conclusion

The Agentic AI Platform for Autonomous End-to-End Revenue Cycle Management represents a paradigm shift in healthcare financial operations.

By combining multi-agent AI orchestration, shared contextual intelligence, and continuous self-learning, the system transforms RCM from a fragmented, manual process into a fully autonomous, adaptive, and high-performance financial engine.

Healthcare organizations adopting this architecture can expect significant gains in efficiency, accuracy, compliance, and revenue realization, positioning them for the next era of intelligent healthcare operations.

Authors:

Anil Reddy

Anil Reddy is SVP- Technology at GeBBS Healthcare Solutions. He is tech enthusiast working on building innovation using AI as a tool and is passionate about bringing in transformation using all forms of technology.

Tarun Agarwal

Tarun Agarwal is Director-AI at GeBBS Healthcare Solutions. A proven leader with significant expertise around Agentic AI, Generative AI, NLP, ML, DL, Python. Helping GeBBS in transforming traditional AI solutions using Agentic/Gen AI.